

ABSTRACT

Title of Document: PHYSICS-BASED DETECTION OF
SUBPIXEL TARGETS IN HYPERSPECTRAL
IMAGERY

Joshua Bret Broadwater
Doctor of Philosophy, 2007

Directed By: Professor Ramalingam Chellappa
Department of Electrical and Computer
Engineering

Hyperspectral imagery provides the ability to detect targets that are smaller than the size of a pixel. They provide this ability by measuring the reflection and absorption of light at different wavelengths creating a spectral signature for each pixel in the image. This spectral signature contains information about the different materials within the pixel; therefore, the challenge in subpixel target detection lies in separating the target's spectral signature from competing background signatures. Most research has approached this problem in a purely statistical manner. Our approach fuses statistical signal processing techniques with the physics of reflectance spectroscopy and radiative transfer theory. Using this approach, we provide novel algorithms for all aspects of subpixel detection from parameter estimation to threshold determination.

Characterization of the target and background spectral signatures is a key part of subpixel detection. We develop an algorithm to generate target signatures based on radiative transfer theory using only the image and a reference signature without the need for calibration, weather information, or source-target-receiver geometries. For background signatures, our work identifies that even slight estimation errors in the number of background signatures can severely degrade detection performance. To this end, we present a new method to estimate the number of background signatures specifically for subpixel target detection.

At the core of the dissertation is the development of two hybrid detectors which fuse spectroscopy with statistical hypothesis testing. Our results show that the hybrid detectors provide improved performance in three different ways: insensitivity to the number of background signatures, improved detection performance, and consistent performance across multiple images leading to improved receiver operating characteristic curves.

Lastly, we present a novel adaptive threshold estimate via extreme value theory. The method can be used on any detector type – not just those that are constant false alarm rate (CFAR) detectors. Even on CFAR detectors our proposed method can estimate thresholds that are better than theoretical predictions due to the inherent mismatch between the CFAR model assumptions and real data. Additionally, our method works in the presence of target detections while still estimating an accurate threshold for a desired false alarm rate.

PHYSICS-BASED DETECTION OF SUBPIXEL TARGETS IN
HYPERSPSPECTRAL IMAGERY

By

Joshua Bret Broadwater

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2007

Advisory Committee:
Professor Ramalingam Chellappa, Chair
Professor Eyad Abed
Professor Larry S. Davis
Professor Adrian Papamarcou
Professor Min Wu

© Copyright by
Joshua Bret Broadwater
2007

Preface

The data used in this dissertation comes from the RDECOM CERDEC Night Vision & Electronic Sensors Directorate (NVESD) of the U.S. Army. The data was collected at significant expense by NVESD and therefore they reserved the right to approve all publications containing their data. Because the NVESD data contains some of the best examples of subpixel target images available, the NVESD imagery is used throughout this dissertation. In order to use their imagery, we had to receive approval from NVESD to publish this dissertation – a ten week process. To help minimize the approval process which dictates that any publication changes must be approved by NVESD, we rewrote the dissertation such that it contains a data chapter. NVESD only requires that this data chapter be approved per e-mail of Mr. David Hicks (NVESD). Fortunately, the addition of this data chapter has provided the added benefit of providing a good explanation of hyperspectral imagery and its idiosyncrasies to motivate the rest of the dissertation.

Dedication

To Andra

Acknowledgements

First, I want to thank my advisor, Prof. Rama Chellappa, for his patience, support, and encouragement while working on this dissertation. He provided me the opportunities to learn, publish, and develop as a scholar and I will always fondly remember my years at the University of Maryland because of those experiences.

Second, I want to thank Dr. Amit Banerjee, Dr. Marc Kolodner, Dr. Reuven Meth, and Dr. Patricia Murphy for their many helpful discussions both on hyperspectral image analysis and their own Ph.D. experiences.

Third, I must thank Ms. Miranda Schatten and Mr. David Hicks of the U.S. Army RDECOM CERDEC NVESD for providing the data used in this dissertation. Hyperspectral data with subpixel targets and detailed ground truth is difficult to find. Without the NVESD data, many of the developments introduced in this dissertation would not have been possible.

Most importantly, I dedicate this dissertation to my wonderful and loving wife, Andra. Without her patience and support, this dissertation would not exist. I also want to thank her for giving me the two best gifts in the world: Sarah and Zachary. I love and appreciate you all more than I can ever express in words.

Table of Contents

PREFACE.....	II
DEDICATION.....	III
ACKNOWLEDGEMENTS.....	IV
TABLE OF CONTENTS.....	V
LIST OF TABLES.....	VII
LIST OF FIGURES.....	VIII
LIST OF ABBREVIATIONS.....	IX
CHAPTER 1: INTRODUCTION.....	1
1.1. A Brief History of Imaging Spectroscopy.....	1
1.2. Subpixel Detection.....	4
1.3. Thesis.....	5
CHAPTER 2: HYPERSPECTRAL DATA.....	10
2.1. AVIRIS.....	10
2.1.1. Sensor Details.....	10
2.1.2. Imagery.....	10
2.2. Sensor X.....	11
2.2.1. Sensor Details.....	11
2.2.2. Imagery.....	12
2.2.3. Spectral Signatures.....	14
2.2.4. Ground Truth.....	16
CHAPTER 3: TARGET SIGNATURE CHARACTERIZATION.....	19
3.1. A Review of Radiometry.....	22
3.1.1. Sun Light.....	24
3.1.2. Sky Light.....	28
3.1.3. Upwelled Radiance.....	30
3.1.4. Atmospheric Transfer Function.....	31
3.2. Current Target Characterization Algorithms.....	31
3.2.1. Model-Based Methods.....	32
3.2.2. In-Scene Methods.....	34
3.3. Average Relative Radiance Transform.....	38
3.4. Experimental Results.....	47
3.4.1. Comparison of Target Radiance Signatures.....	48
3.4.2. Comparison of Target Signatures for Subpixel Detection.....	51
3.5. Summary.....	59
CHAPTER 4: BACKGROUND SIGNATURE CHARACTERIZATION.....	61
4.1. A Review of Endmember Extraction Methods.....	62
4.2. Selected Endmember Extraction Techniques.....	65
4.3. Dimensionality of Hyperspectral Imagery.....	66
4.3.1. Intrinsic Dimensionality Metrics.....	67
4.3.2. Virtual Dimensionality Metrics.....	69
4.4. Experimental Results.....	73
4.4.1. Individual Image Results.....	74
4.4.2. ROC Results.....	78

4.4.3. Conclusions.....	82
4.5. Summary.....	82
CHAPTER 5: PHYSICS-BASED HYBRID DETECTORS.....	84
5.1. Current Subpixel Algorithms.....	87
5.1.1. Fully Constrained Least Squares (FCLS).....	87
5.1.2. Adaptive Matched Subspace Detector (AMSD).....	90
5.1.3. Adaptive Cosine/Coherent Detector.....	92
5.2. Hybrid Detectors.....	95
5.2.1. Hybrid Structured Detector.....	95
5.2.2. Hybrid Unstructured Detector.....	97
5.3. Experimental Results.....	98
5.3.1. Experimental Design.....	99
5.3.2. Endmember Sensitivity Analysis.....	101
5.3.3. Separability Analysis.....	105
5.3.4. Receiver Operating Characteristics.....	110
5.3.5. Conclusions.....	113
5.4. Summary.....	115
CHAPTER 6: ADAPTIVE DETECTION THRESHOLDS VIA EXTREME VALUE THEORY.....	117
6.1. Extreme Value Theory.....	121
6.1.1. The Fisher-Tippett Theorem.....	121
6.1.2. EVT for the Exponential Class.....	122
6.1.3. Generalized Pareto Distribution.....	124
6.2. EVT Adaptive Threshold Algorithm.....	126
6.3. Experimental Results.....	132
6.3.1. Experiments with Known Distributions.....	133
6.3.2. Experiments on Subpixel Target Detectors.....	135
6.3.3. Conclusions.....	142
6.4. Summary.....	143
CHAPTER 7: SUMMARY.....	144
7.1. Cumulative Performance Results.....	144
7.2. Future Work.....	147
7.3. Contributions.....	150
BIBLIOGRAPHY.....	152

List of Tables

Table 1: Description of Sensor X Imagery	14
Table 2: Description of Targets	15
Table 3: Target Ground Truth.....	17
Table 4: Quantitative Comparison of Atmospheric Compensation Algorithms.....	50
Table 5: Comparison of Dimensionality Estimates for Target 1	75
Table 6: Comparison of Dimensionality Estimates for Target 2	75
Table 7: Comparison of Dimensionality Estimates for Target 3	76
Table 8: Comparison of Dimensionality Estimates for Target 4	77
Table 9: Subpixel Experiment Details	99
Table 10: Endmember Sensitivity Results.....	104
Table 11: Comparison of MC and GPD on Known Distributions.....	134
Table 12: Comparison of Threshold Estimates for ACE Results	136
Table 13: Comparison of P_d Estimates for ACE Results.....	137
Table 14: Comparison of False Alarms for ACE Results.....	138
Table 15: Comparison of Threshold Estimates for HSD Results	140
Table 16: Comparison of P_d Estimates for HSD Results.....	141
Table 17: Comparison of False Alarm Rates for HSD Results	142

List of Figures

Figure 1: Hyperspectral Signatures of Common Materials	3
Figure 2: Subpixel Detection Block Diagram.....	7
Figure 3: AVIRIS Image of Cuprite, Nevada	11
Figure 4: Sensor X 1200m Imagery.....	13
Figure 5: Sensor X 300m Imagery.....	14
Figure 6: Target Reflectance Signatures.....	16
Figure 7: Target 3 Radiance Signatures in Image 7.....	18
Figure 8: Target 4 Radiance Signatures in Image 7.....	18
Figure 9: The five sources of light in the reflective wavelengths.....	23
Figure 10: Source-Target-Receiver Geometry.....	25
Figure 11: The Solar Spectrum.....	26
Figure 12: Comparison of Mean Radiance and Reflectance Estimates Using ARRT.....	42
Figure 13: ARRT Block Diagram.....	46
Figure 14: Comparison of Atmospheric Compensation Algorithms for Target 3	49
Figure 15: Comparison of Atmospheric Compensation Algorithms for Target 4	49
Figure 16: ACE Results for Image 7.....	54
Figure 17: ROC Comparison of Target 1 Signatures.....	57
Figure 18: ROC Comparison of Target 2 Signatures.....	57
Figure 19: ROC Comparison of Target 3 Signatures.....	58
Figure 20: ROC Comparison of Target 4 Signatures.....	59
Figure 21: Comparison of Background Dimension Estimates for Target 1	79
Figure 22: Comparison of Background Dimension Estimates for Target 2	80
Figure 23: Comparison of Background Dimension Estimates for Target 3	81
Figure 24: Comparison of Background Dimension Estimates for Target 4	81
Figure 25: Graphical Comparison of Endmember Sensitivity.....	103
Figure 26: Separability Analysis for Target 1.....	106
Figure 27: Separability Analysis for Target 2.....	107
Figure 28: Separability Analysis for Target 3.....	108
Figure 29: Separability Analysis for Target 4.....	109
Figure 30: Subpixel Detection ROC Curves for Target 1.....	111
Figure 31: Subpixel Detection ROC Curves for Target 2.....	111
Figure 32: Subpixel Detection ROC Curves for Target 3.....	112
Figure 33: Subpixel Detection ROC Curves for Target 4.....	112
Figure 34: Comparison of the GPD to the Empirical CDF for Example 1.....	130
Figure 35: Comparison of the GPD to the Empirical CDF for Example 2.....	131
Figure 36: Comparison of Corrected Samples.....	131
Figure 37: Block Diagram of the EVT Adaptive Threshold Algorithm.....	132
Figure 38: Proposed Subpixel Detection Block Diagram.....	145
Figure 39: Subpixel Detection System ROC Curves.....	146

List of Abbreviations

ACE.....	Adaptive Coherent/Cosine Estimate
AIC.....	Akaike Information Criterion
AMEE	Automated Morphological Endmember Extraction
AMSD	Adaptive Matched Subspace Detector
ARRT	Average Relative Radiance Transform
AVIRIS	Airborne Visible Infrared Imaging Spectrometer
BIC.....	Bayesian Information Criterion
BRDF	Bidirectional Reflectance Distribution Function
CDF.....	Cumulative Distribution Function
CEM.....	Constrained Energy Minimization
CFAR	Constant False Alarm Rate
CSD.....	Constrained Signal Detector
EIF.....	Empirical Indicator Function
ELM	Empirical Line Method
EVT.....	Extreme Value Theory
FCLS	Fully Constrained Least Squares
GEVT	Generalized Extreme Value Theory
GLRT	Generalized Likelihood Ratio Test
GPD.....	Generalized Pareto Distribution
GPS	Global Positioning System
HIS	Hyperspectral Imagery
HSD.....	Hybrid Structured Detector
HUD.....	Hybrid Unstructured Detector
IARR	Internal Average Relative Reflectance
IEA.....	Iterative Error Analysis
IS.....	Importance Sampling
JHU/APL.....	The Johns Hopkins University Applied Physics Laboratory
JPL	Jet Propulsion Laboratory
LSE	Least Squares Error
LVQ	Learning Vector Quantization
LWIR	Long Wave Infrared
MC.....	Monte Carlo
MDL.....	Minimum Description Length
MEI	Morphological Eccentricity Index
MLE	Maximum Likelihood Estimate
MNF.....	Maximum Noise Fraction
MODTRAN	Moderate Transmission model
MSMA	Modified Spectral Mixture Analysis
MWIR	Mid-Wave Infrared
NASA.....	National Aeronautics and Space Administration
NIR.....	Near Infrared
NRL.....	United States Naval Research Laboratory
NSP	Noise Subspace Projection

NVESD	Night Vision & Electronic Sensors Directorate
ORASIS	Optical Real-Time Spectral Identification System
OSP	Orthogonal Subspace Projection
PCA	Principal Component Analysis
PDF	Probability Density Function
PPI	Pixel Purity Index
ROC	Receiver Operating Characteristic
SAA	Simulated Annealing Algorithm
SAM	Spectral Angle Mapper
SINR	Signal to Interference plus Noise Ratio
SOLCD	Spectral Object Level Change Detection
SMM	Stochastic Mixing Model
SVD	Singular Value Decomposition
SWIR	Short Wave Infrared
USGS	United States Geological Survey
VIS	Visible

Chapter 1: Introduction

1.1. A Brief History of Imaging Spectroscopy

The study of a material's spectral properties grew out of the field of reflectance spectroscopy introduced in the 1920s. Reflectance spectroscopy identified the component chemicals in a sample by studying the reflective properties of the material [40]. By the 1930s and 1940s, spectrophotometers were introduced and the field of spectroscopy grew more popular. This work led to radiative transfer theory that was able to measure the reflective properties of a sample and identify the underlying physical mechanisms in such measurements. Radiative transfer theory ultimately led to the development of spectral imagers in the early 1970s [54].

Spectral imagery is, however, not a new concept. Color imagery is the most basic and widely recognized spectral imagery. In spectral imagery, each spatial point or pixel is represented by multiple measurements of different wavelengths in the electromagnetic spectrum. In the case of color imagery, each pixel contains information for the red, green, and blue wavelengths in the visible portion of the electromagnetic spectrum. This idea of measuring the energy in different wavelengths of the spectrum along with radiative transfer theory led to the development of multispectral imagery.

In July 1972, the first space-based multispectral imager was launched under the LANDSAT program [63]. The imager contained four bands across the visible (VIS) to near-infrared (NIR) wavelengths. The LANDSAT program was so successful that the program continues today utilizing new multispectral sensors that are capable of measuring seven bands of the electromagnetic spectrum. The success

of these multispectral sensors led to the development of the hyperspectral sensor in the mid-1980s and its corresponding field of imaging spectroscopy.

Hyperspectral imagery (HSI) differs from its earlier counterpart, multispectral imagery, in two key ways. The first difference is the number of spectral bands collected by hyperspectral sensors. Multispectral sensors typically collect less than ten bands of spectral information per pixel. Hyperspectral imagery contains hundreds of bands of spectral information per pixel. The second difference is that multispectral imagery having so few bands, selects wavelengths that are considered the most informative for a particular application; thus, the bands are non-contiguous. Hyperspectral sensors sample the spectrum creating hundreds of contiguous spectral bands. The result is a spectral signature at every pixel location that can be used to identify the materials imaged within the pixel. The spectral signature can also be decomposed to identify different materials present in the same pixel.

For this dissertation, we focus on hyperspectral sensors that measure energy in the reflectance wavelengths of the electromagnetic spectrum. Reflectance is defined as “the ratio of reflected radiance to incident irradiance” [93]. Simply, reflectance is a measure of the energy reflected from the surface of an object. Therefore, hyperspectral sensors in the reflective wavelengths are passive instruments measuring the light reflected in a scene – typically sunlight. The reflectance wavelengths in the electromagnetic spectrum are composed of three spectral bands: the Visible (VIS) from 400 nm to 700 nm, the Near Infrared (NIR) from 700 nm to 1100 nm, and the Short Wave Infrared (SWIR) from 1100 nm to 2500 nm. Figure 1 displays these three spectral bands and provides three typical materials in a hyperspectral image: road,

soil, and vegetation. This shows figure demonstrates the spectral resolution available in hyperspectral imagery.

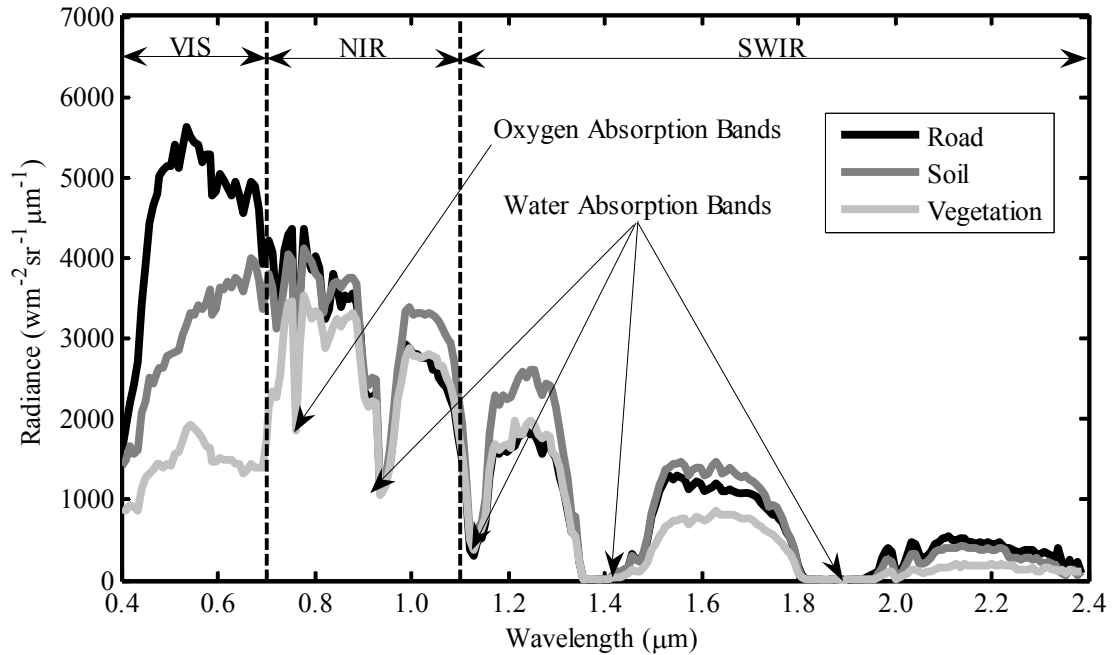


Figure 1: Hyperspectral Signatures of Common Materials

Figure 1 also displays a few of the effects caused by light passing through the atmosphere. Therefore, hyperspectral sensors do not directly measure the reflectance properties of a material. Instead, hyperspectral sensors measure the radiance at each wavelength. Radiance is defined as “radiant flux per unit area per unit solid angle per unit wavelength” [93]. The radiance values not only contain the reflectance properties of the object being imaged, but also contain all of the environmental effects that arise between the imager and the object being imaged. Thus, the hyperspectral sensor not only records the materials in the pixel, but also the spectral signatures due to sunlight and the atmosphere such as the absorption bands shown in Figure 1.

Despite the effects of the atmosphere masking the true reflective signatures of the materials being imaged, a number of applications have been developed to utilize

hyperspectral imagery such as mineral identification [76][77], land cover classification [34], vegetation studies [66], and atmospheric studies [72]. This dissertation focuses on target detection applications – specifically, subpixel detection where the target is literally smaller than the area imaged by a single pixel. This field of study has broad reaching applications from obvious military applications to search and rescue operations [106] to forensic investigations for the space shuttle Columbia incident [78]. The last application is perhaps the most well known use of hyperspectral sensors to perform broad-area searches and find parts of the Columbia that were only one inch long from an altitude of 2000 ft.

1.2. Subpixel Detection

Detection can be considered a special two class case of pattern recognition; however, it differs from classification in a number of ways [69]. In classification, the objective is to minimize the total error across all classes of data [24]. In detection, we only want to identify our desired target class amongst a larger background class. This reasoning fundamentally assumes that the target class is rare and that most pixels are from the background class. Thus, if we minimized the total error as in classification, we could simply identify every pixel as background. Of course, we are interested in maximizing the detection of targets while minimizing Type I errors – identifying background pixels as targets (false alarms) [18]. This maximization of target detection and minimization of false alarms is the fundamental difference between detection and standard pattern recognition.

Spectral subpixel detection in hyperspectral image (HSI) data aims to identify a target smaller than the size of a pixel using only spectral information [71]. Thus, the challenge in detecting subpixel targets lies in separating the target's spectral signature

from other competing signatures within the pixel. To accomplish this “unmixing” of signatures, the field of reflectance spectroscopy provides a model of how these multiple spectra interact with one another [40]. The most common model assumes that the spectra are represented by unique spatially non-overlapping materials. This model is called the linear mixing model and it is the cornerstone for most subpixel detection algorithms.

The linear mixing model assumes that a pixel is made up of endmembers, each with its own abundance. Endmembers are the spectra representing the unique materials in a given image. For instance, in an image that contains soil, vegetation, and road, the endmembers would be the corresponding unique spectral signatures for each of these materials as shown in Figure 1. Abundances are the percentage of each material within a given pixel. Mathematically, the linear mixing model is written as

$$\mathbf{x} = \mathbf{E}\mathbf{a}, \quad a_i \geq 0, \quad \sum_{i=1}^M a_i = 1 \quad (1)$$

where \mathbf{x} is an $L \times 1$ vector that represents the spectral signature of the current pixel, M is the number of endmembers within the image, \mathbf{E} is an $L \times M$ matrix where each column represents the i^{th} endmember, and \mathbf{a} is an $M \times 1$ vector where the i^{th} entry represents the abundance value a_i . Note that the linear mixing model includes two constraints on the abundance values: non-negativity and sum-to-one. These constraints place physical limitations on the abundances making sure they represent the percentage of each material present in the pixel.

1.3. Thesis

The interesting part of subpixel detection is not the linear mixing model itself, but the parameters of the linear mixing model. These parameters have been

historically treated only in a statistical sense. The parameters are typically found using maximum likelihood estimates (MLE). This is, of course, a natural way to proceed in solving detection problems since such estimates are guaranteed to be consistent and asymptotically efficient [18]. However, Prof. David Landgrebe, a pioneer in remote sensing, argues in his paper that the improvement in hyperspectral image analysis will not be made by using different statistical algorithms, but by properly modeling the physics of the problem [64]. Instead of using statistical estimates of the parameters, we could use physics-based estimates of the parameters within statistical hypothesis tests to improve subpixel detection.

Some research has already been devoted to this type of physics-based detection approach. The most notable is from Thai and Healey [109]. They present an algorithm that creates a subpixel detector that is invariant to atmospheric effects. They project the desired target reflectance signature to radiance signatures for thousands of different atmospheric profiles using the computational physics model MODTRAN (MODerate TRANsmission) [3]. From these thousands of possible target radiance signatures, they use singular value decomposition (SVD) to extract a set of target singular vectors that minimize atmospheric and illumination effects; however, they only use physics to derive the target signature. The background signatures and detector are still estimated using purely statistical arguments. This has the negative effect of generating abundances that cannot meet the linear mixing model constraints.

Schott [94] and Lee [65] take a slightly different approach to physics-based subpixel detection. From the thousands of different target radiance signatures generated with MODTRAN, Lee uses a simplex method to identify the target

signatures that span the space of all possible target signatures generated. These target “endmembers” are concatenated to the image data and a simplex method such as N-FINDR is used to extract the endmembers [115][116] – some of which they argue will be target signatures. This has the result of creating both target and background endmembers that are physically meaningful. Unfortunately, they too use least squares estimates of the abundances even though physically meaningful abundances could be estimated from their endmember signatures.

Our physics-based subpixel detection approach uses physically meaningful estimates of both the endmembers and their abundances. We show this approach leads to not only improved detection performance over previous approaches, but also provides a level of insensitivity to estimation errors and provides contextual information not obtainable with other methods. Additionally, we propose new algorithms for nearly all facets of subpixel detection (shown in Figure 2) from parameter characterization to threshold estimation.

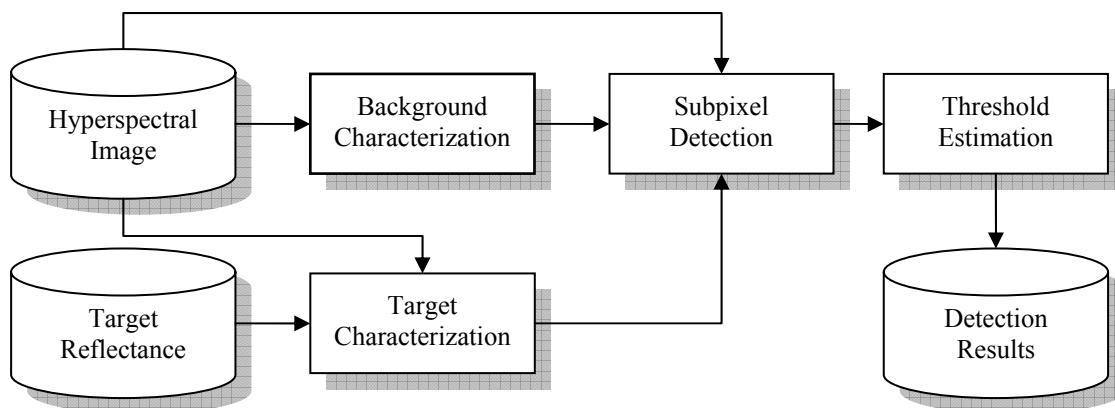


Figure 2: Subpixel Detection Block Diagram

In Chapter 3, we present a novel way to estimate target radiance signatures from reflectance measurements using only the target reflectance signature and the

hyperspectral image. This chapter provides an overview of radiative transfer theory and how MODTRAN and other methods use this theory to estimate radiance signatures from reflectance measurements. We explain how MODTRAN can be used with proper weather, topographic, and geometric data to generate a target signature for a specific hyperspectral image. From this, we develop a new in-scene algorithm that performs similarly to MODTRAN, but uses only a target and reference reflectance signature along with the hyperspectral image to estimate a target radiance signature for subpixel detection.

In Chapter 4, we present a new method to estimate the number of endmembers that maximize subpixel detection performance. The chapter gives a brief overview of endmember extraction techniques and identifies the algorithms we use in this dissertation to obtain physically meaningful endmembers. The chapter documents the sensitivity of subpixel target detection to the number of endmembers showing how slight errors in estimating the number of endmembers can cause severe losses in performance. From this result, we compare a number of different algorithms to estimate the number of endmembers and compare them to our proposed methods relative to subpixel detection performance.

In Chapter 5, we present our physics-based hybrid subpixel detectors [12]. Unlike the subpixel detectors proposed by [41], [49], [58], and [71], we develop a detector that uses all of the linear mixing model constraints including the non-negativity and sum-to-one constraints of the abundances. Our work differs from previous work because of how it models the data. The assumption in the literature is that the error between the linear mixing model and HSI data can be modeled by zero-

mean noise with a covariance matrix of $\sigma^2 \mathbf{I}$. This has been shown to be erroneous in [71]. Using this result, we model the remaining noise using a full covariance matrix to account for sensor artifacts and nonlinear mixing effects not represented by the linear mixing model. This results in a subpixel detector that has improved performance and is partially insensitive to the number of background endmembers used.

In Chapter 6, we present a new algorithm to estimate a detection threshold for a desired false alarm rate for any detector. One of the disadvantages of the hybrid subpixel detectors is the use of the non-negativity constraints of the linear mixing model. These constraints disallow a closed-form solution for the detector making derivation of the target and background conditional distributions difficult at best. To overcome this shortfall, we develop an adaptive threshold technique based on Extreme Value Theory (EVT). We show the proposed technique outperforms both theoretical estimates for Constant False Alarm Rate (CFAR) detectors as well as non-parametric methods such as Monte Carlo estimates – especially when targets are present in the imagery.

In Chapter 7, we summarize our work and present an example of the proposed algorithms working together in a subpixel detection process. Besides providing excellent detection of subpixel targets, the result shows the ability of these methods to provide near real-time results using a minimal amount of ancillary information. This result is important to transitioning hyperspectral subpixel detection algorithms from research to practice.

Chapter 2: Hyperspectral Data

In this dissertation, we use hyperspectral imagery from two sensors: the Airborne Visible Infrared Imaging Spectrometer (AVIRIS) and the U.S. Army RDECOM CERDEC Night Vision & Electronic Sensors Directorate (NVESD) Sensor X. The chapter is therefore broken into two sections. Each section contains information about the hyperspectral sensor, its images, available target reflectance signatures, and corresponding ground truth information.

2.1. *AVIRIS*

2.1.1. *Sensor Details*

The AVIRIS imagery comes from the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL) at the California Institute of Technology [111]. This sensor collects 224 contiguous spectral bands spanning the wavelengths from 400 to 2500 nm. The sensor was primarily designed for environmental remote sensing applications; therefore, the imagery collected has not been focused on subpixel detection applications. Nevertheless, the AVIRIS sensor has been well calibrated and does not contain any low SNR bands allowing us to use all 224 spectral bands for processing.

2.1.2. *Imagery*

We chose one image to use from the AVIRIS data sets: the Cuprite, Nevada image [107]. From the Cuprite data set, we chose a sub-image containing a small town shown in Figure 3. The image itself covers a 10.4 km by 5.1 km swath of area with each pixel measuring 17 m per side. While the AVIRIS imagery has not been focused on subpixel detection applications, it can be useful to demonstrate the

atmospheric compensation techniques in Chapter 3. AVIRIS images are delivered as two images: the original radiance image collected by the sensor and another image which is an estimate of the reflectance signatures at each pixel in the image using known ground materials. These reflectance estimates will be used to identify how well our proposed target characterization method identifies radiance signatures generated from flat reflectance signatures.

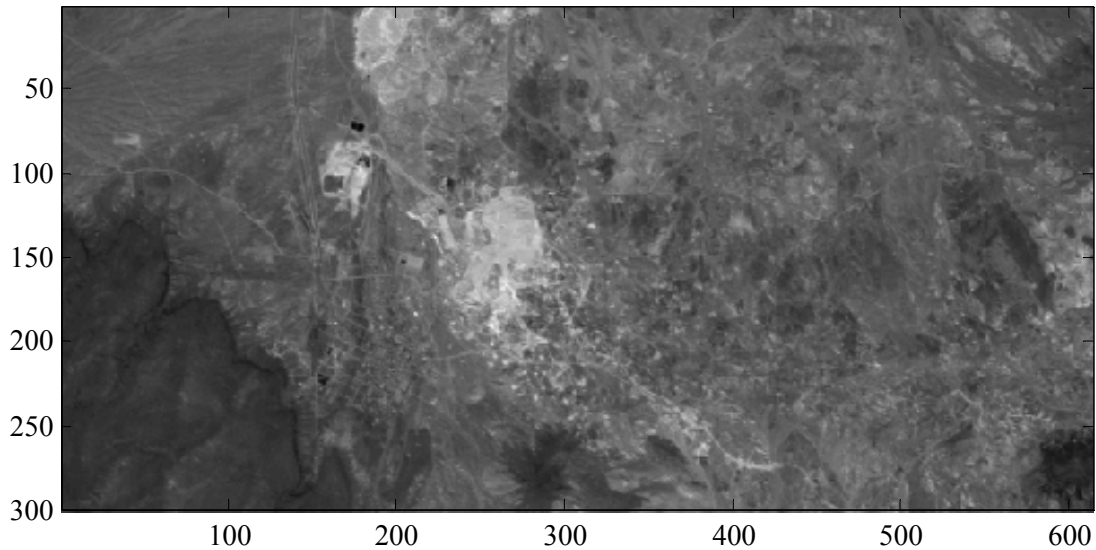


Figure 3: AVIRIS Image of Cuprite, Nevada

2.2. *Sensor X*

2.2.1. *Sensor Details*

The Sensor X imagery comes from the U. S. Army RDECOM CERDEC Night Vision & Electronic Sensors Directorate (NVESD). The sensor collects 256 contiguous spectral bands spanning the wavelengths from 400 to 2500 nm. Along with the sensor specifications, we received a spreadsheet containing information about the sensor's spectral bands. For example, the absorption bands for oxygen, carbon dioxide, and water were well documented. The spreadsheet also identified low

SNR bands in the imagery due to sensor artifacts. For our target detection application, these bands are non-informative and only serve to increase processing time without providing any benefits. Because of this, we did not use these bands as is typically done in target detection applications [41],[70],[71]. After removing these bands, we are left with 169 spectral bands for our subpixel detection experiments.

2.2.2. Imagery

We chose seven images to use in this dissertation. The first six images were chosen because of their small fill factors (e.g., percentage of a pixel that is comprised of target) and the difficult background in which the targets lie. The most difficult of these areas is the tall grass site. At this site, the grass is high enough to partially obscure the target causing the pixel fill factors to be smaller than expected. The other two areas are easier since the targets are not obscured. Figure 4 shows the six images with corresponding target locations.

The seventh image is shown in Figure 5. This image was chosen because the targets were full or multi-pixel. This image was selected because the true target radiance signatures could be extracted from the image. These signatures can be compared to the target radiance estimates described in Chapter 3.. Without this image, we would not know how well the target characterization algorithms were performing. The image is only used for Chapter 3. Table 1 identifies each of the images, the type of area imaged, the amount of area imaged, and the spatial resolution of an individual pixel.

Unfortunately, the imagery we received was collected with an uncalibrated sensor. This posed a significant problem. Some of the algorithms within this dissertation use the physics-based model MODTRAN that calculates the radiance of

an object from its corresponding reflectance signature. The radiance signature generated by the model assumes the sensor is calibrated. When the sensor is not calibrated, the model will predict signatures that will not match those in the imagery. This mismatch is severe enough to render a target detection algorithm useless.

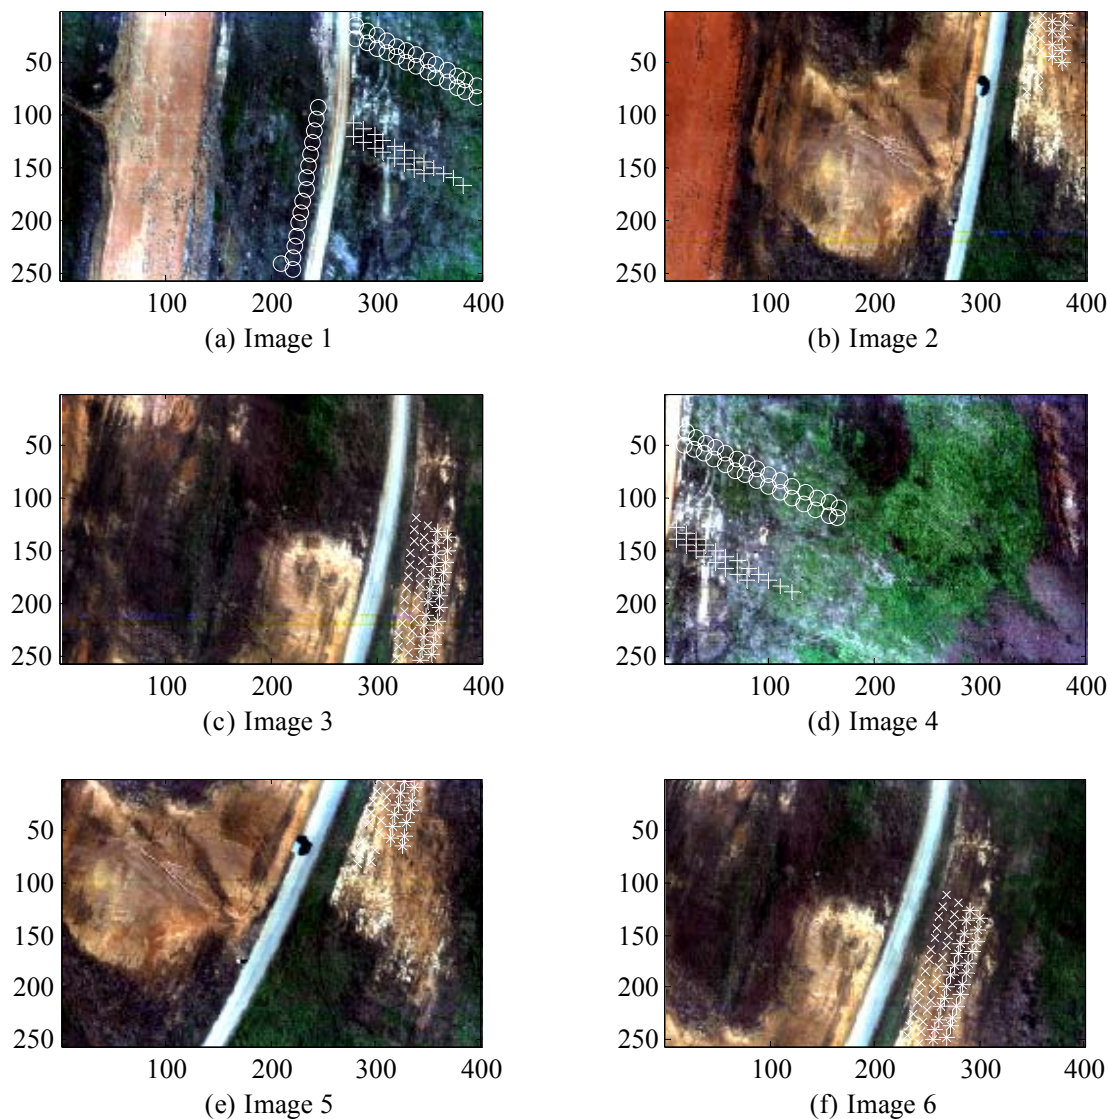


Figure 4: Sensor X 1200m Imagery
(Target 1 '+', Target 2, 'o', Target 3 'x', Target 4 '*')

To overcome this problem, we worked with Dr. Marc Kolodner of the Johns Hopkins University Applied Physics Laboratory (JHU/APL). Using MODTRAN, we generated radiance signatures for known background materials in the imagery. We

compared the model-based signatures to the known signatures in the imagery. From these comparisons, an offset and gain vector was created. This offset and gain was applied to each image to vicariously calibrate the image. These new vicariously calibrated images were then used for the experiments in this dissertation.

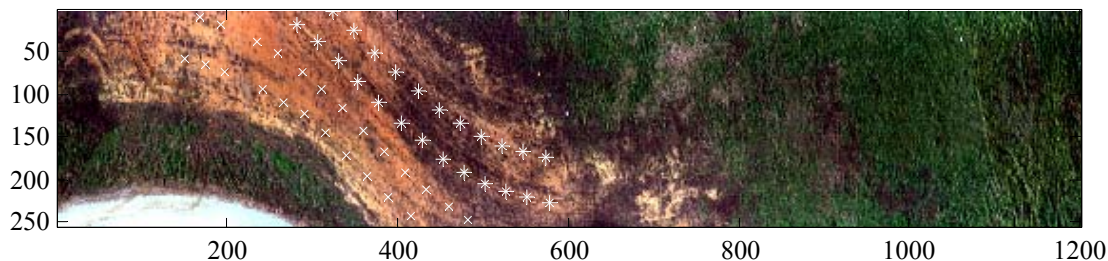


Figure 5: Sensor X 300m Imagery
(Target 3 'x', Target 4 '*')

Table 1: Description of Sensor X Imagery

Image	Background	Clutter Density	Altitude (m)	Area (m ²)	Pixel Size (m ²)
1	Short & Tall Grass	High	1220	18811	0.1823
2	Sparse Grass	Medium	1220	18811	0.1823
3	Sparse Grass	Medium	1220	19464	0.1823
4	Short Grass	Medium	1216	18815	0.1815
5	Sparse Grass	Medium	1215	18542	0.1806
6	Sparse Grass	Medium	1213	19097	0.1806
7	Sparse Grass	Medium	313	7400	0.0241

2.2.3. Spectral Signatures

Besides the imagery, we received spectral libraries containing reflectance signatures for both the targets and background materials. All signatures were collected using hand-held spectrometers in the field. Due to this in-field data capture, multiple signatures were created for each target and background material. These signatures were averaged to form a signature for each material. This method was chosen because the averaged spectral signature reduced variations that occurred when measuring with the hand-held spectrometer.

For the background, numerous signatures were collected. These ranged from different types of vegetation to fiducial markers placed in the field for spatial registration purposes. This information is typically not available in real-world applications, but allows us to vicariously calibrate the images. The signatures are also used as reference signatures to help estimate the amplitude of the target signature as explained in Chapter 3.

From the target signatures, we chose four different targets. The targets were chosen to provide a wide variety of spectral signatures. The targets are typically pieces of metal or plastic small enough to achieve subpixel sizes at 1200m altitudes. Additionally, the targets have different paints which cause the reflectance signatures to vary from very strong (Target 1) to very weak (Target 4) as shown in Figure 6. Table 2 provides a description of each target's geometry, size, material, color, and symbol used in figures throughout the dissertation.

Table 2: Description of Targets

Target	Geometry	Size (m ²)	Material	Color	Symbol
1	Circle	0.0182	Plastic	White	+
2	Circle	0.0869	Metal	Green	o
3	Square	0.1090	Plastic	Green	x
4	Circle	0.0869	Metal	Dark Green	*

Target 3 was an interesting case as that particular target had two spectral signatures. The two signatures existed because it was discovered later that the targets were made of slightly different plastics. The difference was very slight as can be seen in Figure 6, but was significant enough that it was decided two signatures should be used. We chose to use this target because it is the only case where we have multiple target signatures for a single target type.

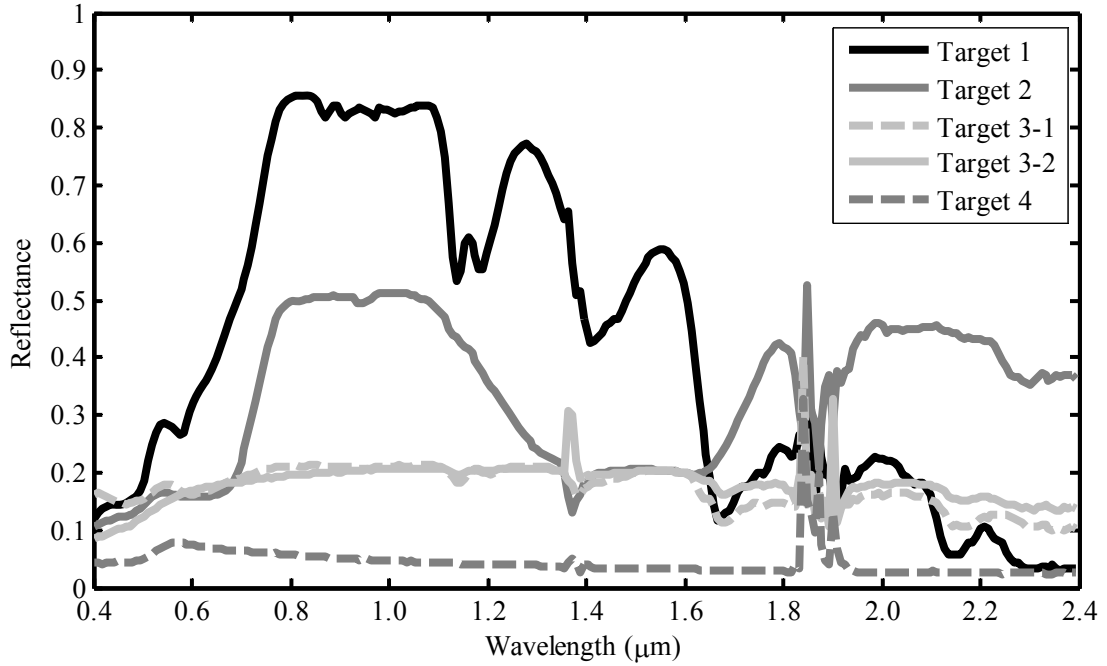


Figure 6: Target Reflectance Signatures

2.2.4. Ground Truth

Along with the imagery and signatures we received from NVESD, we received ground truth information identifying the target locations in the imagery. The ground truth data contained object-level location information. Unlike pixel-level truth which identifies the location of the targets for each pixel and their corresponding abundances, object-level truth specifies an area in the image where the targets are located. Therefore, the ground truth identifies the center of the target even though it may span multiple pixels. Note that this statement is true even with subpixel targets as the target could be located on pixel borders. Table 3 details how many targets are in the seven images arranged by target type and image. The locations of each target in the Sensor X imagery can be seen in Figure 4 and Figure 5.

Given object-level ground truth, we had to cluster the detector outputs to form objects as pixel level analysis was not possible. To obtain these objects, a clustering

threshold is applied to each image. This clustering threshold refers to a threshold that combines adjacent pixels together to form an object which will be classified as either target or clutter. Typically this threshold is chosen to include no more than 1% to 5% of the pixels in the image depending on the application. In our analysis, we chose 1% as we knew the number of targets was far less than 1% of the pixels in any one image. Each cluster is assigned the maximum detection score from all the pixels that make up the cluster. Along with the maximum detection score, each cluster is identified as either target or clutter based on their location relative to the object-level ground truth. This information can then be used to identify how well a detector performs.

Table 3: Target Ground Truth

Image	Target 1	Target 2	Target 3	Target 4	All
1	20	42	0	0	62
2	0	0	12	9	21
3	0	0	25	23	48
4	20	30	0	0	50
5	0	0	15	12	27
6	0	0	28	25	53
7	0	0	24	24	48
All	40	72	104	93	309

From the ground truth information, we were able to extract target radiance signatures from Image 7 due to the targets spanning multiple pixels. These “true” target radiance signatures will be used in Chapter 3 to compare the estimated target radiance signatures with the ones shown in Figure 7 and Figure 8. Each figure contains all of the target radiance signatures found in the image (in gray) and their spectral average (in black). Note the wide variability of target signatures in either case. Despite our best efforts, some background signatures leaked into our “true” target signatures. This occurred because even with four pixels on target, some small

amounts of background signatures may still be present. This is especially the case for Target 4 where the targets spanned on average 3.6 pixels.

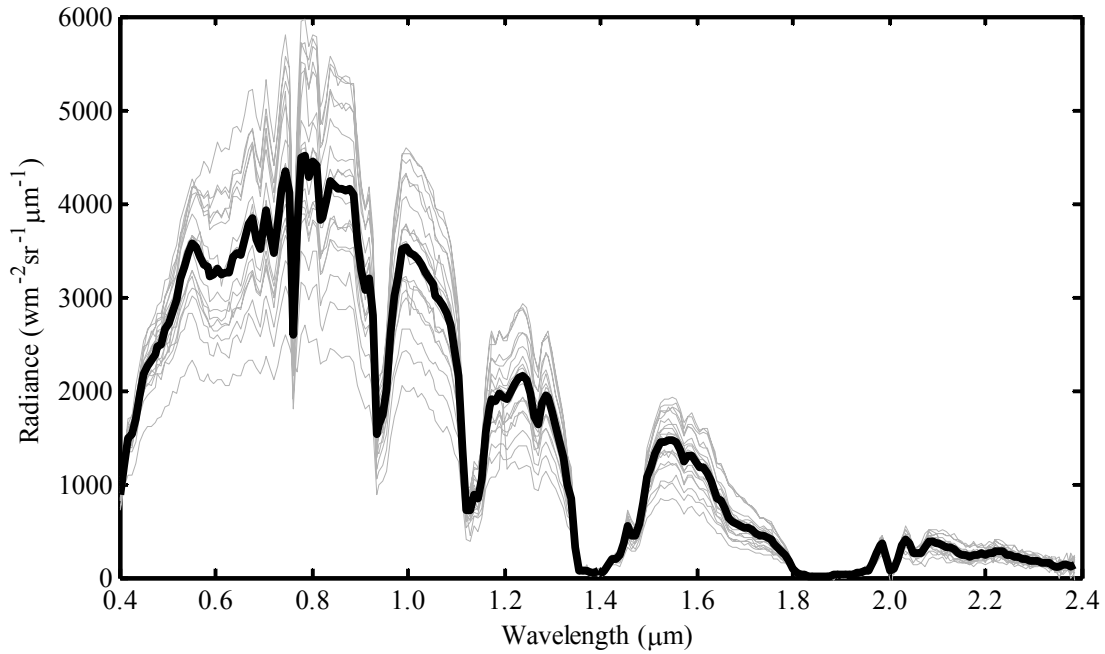


Figure 7: Target 3 Radiance Signatures in Image 7
(Gray lines represent individual targets and black line represents the mean)

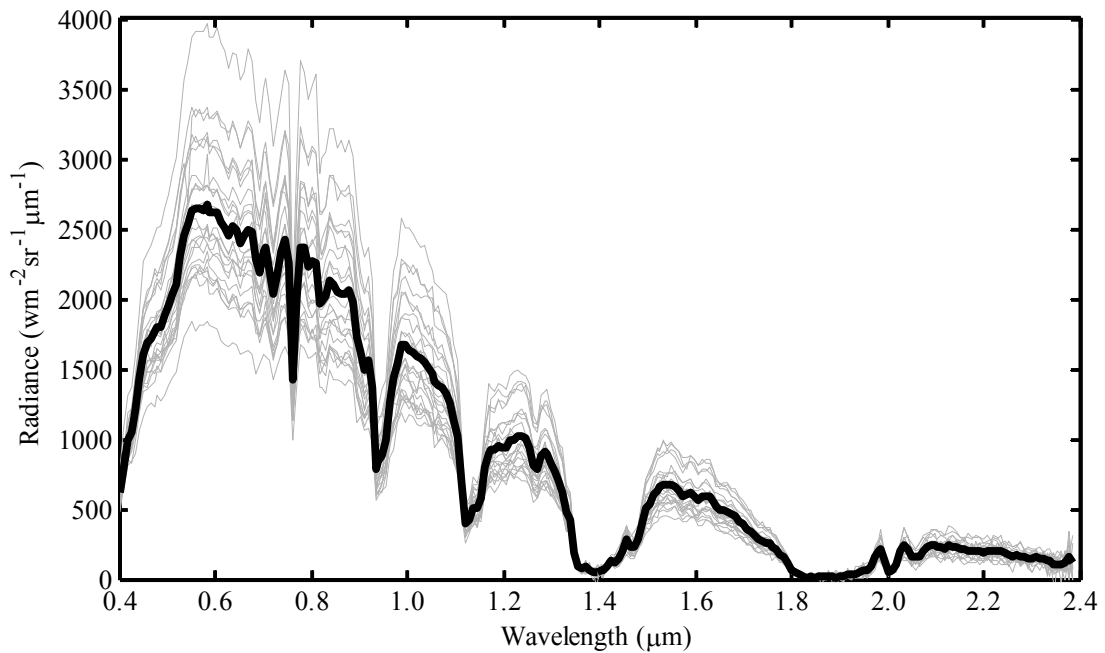


Figure 8: Target 4 Radiance Signatures in Image 7
(Gray lines represent individual targets and black line represents the mean)

Chapter 3: Target Signature Characterization

An important part of subpixel detection is the correct characterization of the target signature. As explained in Chapter 1, target characterization is especially important for hyperspectral detection because the images are collected in terms of radiance while the target signatures are measured in terms of reflectance. The reason for this mismatch is due to the fact that target signatures are typically measured in laboratories or in the field with hand-held spectrometers that are at most a few inches from the target surface. Hyperspectral images, however, are collected hundred to thousands of meters away from the target and have significant atmospheric effects present. Therefore, a transfer function between radiance and reflectance must be obtained. This transfer function is known as atmospheric compensation.

A number of algorithms have been developed to compensate for atmospheric effects. The algorithms can be classified into two primary types: radiance inversion methods and radiance projection methods. Radiance inversion methods were first developed for spectral analysis purposes. Originally, hyperspectral imagery was used to classify images into different natural phenomenon for applications such as mineral mapping [59],[98],[107]. In order to accomplish this type of classification, the logical path was to invert the image from radiance to reflectance and compare the resulting corrected image to known spectral reflectance libraries. The idea in these programs was not to identify a certain material, but to identify the constituent materials in the image for mapping purposes. One such algorithm is FLAASH [3].

While this may be ideal for image analysts wanting to investigate spectral signatures, it is not the best method for detecting subpixel targets. First, the

algorithms process every pixel in the image requiring significant processing time. Second, the algorithms have to make simplifying assumptions to perform the inversion because it is intrinsically an ill-posed problem [75]. So, while these programs have enjoyed some success in target detection applications, they are better suited for spectral analysis by operators that can make informed judgments.

The other class of atmospheric compensation algorithms is based on radiance projection methods. These methods project a reflectance signature into a radiance signature for a particular hyperspectral image. Murphy and Kolodner have one of the most direct approaches: calculate the radiance of a target signature at the sensor using real-time weather predictions and the known source-target-receiver geometry [75]. This type of atmospheric compensation algorithm makes good use of computational physics using the MODTRAN atmospheric model [3]. It also provides different shading conditions so targets can be modeled in both full sun and full shade (such as in the shade of a tree or cloud). Although this approach is the most direct and computationally simple, it also requires the most ancillary information to work properly. Weather data must be timely and the source-target-receiver geometry known precisely. For new data collections, this is usually not hard information to obtain; however, for past data collections, this method typically cannot be used

Healey and Slater simultaneously developed another forward projection model that was designed to be atmospheric invariant [45]. Based on Healey's earlier work with color imagery, they developed an algorithm that projected a target reflectance signature into approximately 17,000 different environments. From these 17,000 radiance signatures, they used SVD to create a nine-dimensional subspace that could

be used in any environment. Results show that this method works well, but requires a significant amount of pre-processing to create the invariant subspace.

A final set of methods use in-scene information to calculate the target radiance signature. These approaches directly estimate atmospheric effects by using information present in the imagery. The most popular of these is the Empirical Line Method [26]. This method uses an adaptive background estimator to find any vegetation in the imagery. Vegetation is used because it is typically ubiquitous and has a well-known reflectance signature. Using the estimated vegetation signature from the image and the known vegetation reflectance signature allows a direct calculation of the transfer function without MODTRAN or any other physical modeling technique. The only issue with such an approach is that certain environments may not have vegetation in the image such as urban environments, winter scenes, or desert scenes.

This chapter presents our work and analysis of model-based and in-scene based radiance projection methods. To begin, we describe in some detail the atmospheric transfer function and the simplifying assumptions made for estimation purposes. We next describe two current methods for atmospheric compensation: an in-scene method developed by Piech and Walker [80] and a model-based method using MODTRAN with radiosonde information. . We then present our own in-scene method for target characterization called Average Relative Radiance Transform (ARRT). The final sections of the chapter compare ARRT to MODTRAN. It will be these two methods which we will use throughout the dissertation for target signature characterization.

3.1. A Review of Radiometry

Radiometry is the measurement of electromagnetic fields typically in the visible and infrared wavelengths [93]. To understand the measurements at an optical sensor, radiometry (or radiative transfer theory) has produced a model of how photons (light) propagate from the sun and through the atmosphere. By understanding this model, we can understand which parts of the radiance signature measured at the sensor are produced by the target of interest and which are produced by the surrounding environment. We can also understand which parts of the model are more critical than others for target characterization.

For this dissertation, we only cover the most basic radiometric principles; however, there are two excellent books available by Schott [93] and Hapke [40] that provide greater details about this interesting theory. Schott's book is meant primarily for the general scientist and engineer interested in remote sensing. Hapke's book provides a more thorough analysis of the governing equations of light. Both are excellent resources and much of the material in this section is derived from both of these texts.

For this dissertation, we are concerned only with those photons that can be collected by a hyperspectral sensor in the reflectance domain. The reflectance domain identifies a range of electromagnetic wavelengths from 400 nm to 2500 nm where light is primarily reflected from objects. As the wavelengths increase, the dominant effect becomes self emittance of photons (such as heat). While this is an interesting regime, our data is all collected in the reflectance wavelengths and as such, we will restrict our analysis to these wavelengths.

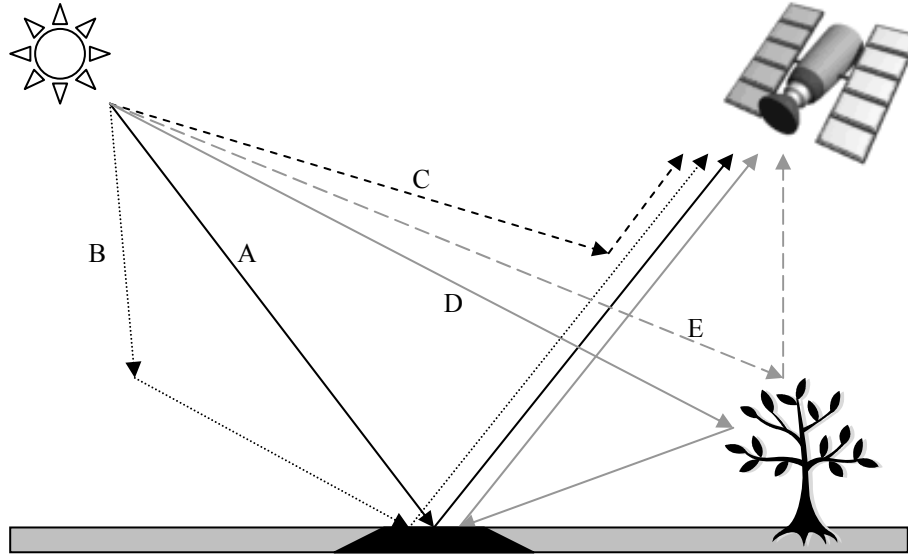


Figure 9: The five sources of light in the reflective wavelengths (A: Direct Sunlight, B: Sky Light, C: Upwelled Radiance, D: Multipath Effect, E: Adjacency Effect)

In the reflectance domain, there are five main sources of light collected by a sensor: direct sun light, sky light, upwelled radiance, multipath effect, and the adjacency effect. These multiple sources of light are shown in Figure 9. Sun light is the light generated by the sun that passes through the atmosphere, reflects off the area being imaged, and is collected at the sensor. Sky light is the light that is scattered in the atmosphere which reflects off the area being imaged and back to the sensor. Upwelled radiance is the light that is scattered in the atmosphere that never reaches the area being imaged. Instead, this light is scattered directly into the optical path of the sensor. Multipath effects are due to light that reflects off of multiple objects in a scene before arriving at the sensor. The adjacency effect occurs when light scatters off of other background objects near the area being imaged into the optical path of the sensor [52]. The last two sources of light are very small compared to the first three

and are typically not computed in most models. Because of these reasons, only the first three light sources will be treated in greater detail.

3.1.1. Sun Light

The most obvious source of light is the sun. Photons are generated at the sun and pass through the atmosphere onto the object being imaged and back to the sensor. Along the way, the spectral properties of the light are changed as the photons are absorbed and scattered through the atmosphere. These effects can be mathematically modeled as

$$L_{sun}(x, y, \lambda) = KT_u(z_g, z_u, \theta_v, \phi_v, \lambda)R(x, y, \lambda)T_d(z_g, \theta_0, \phi_0, \lambda)E_0(\lambda)\cos \theta_0 \quad (1)$$

where L_{sun} is the radiance seen at the sensor generated from sun light, K is the amount of energy at the top of the atmosphere, T_u is the upward atmospheric transmittance, R is the reflectance of the object being imaged, T_d is the downward atmospheric transmittance, and E_0 is the exoatmospheric spectral signature of the sun. All of these quantities are a function of the spectral wavelength λ and most of the quantities are based on the geometry of the source (sun), target (object being imaged), and receiver (camera) geometry as shown in Figure 10. The geometries are based on cylindrical coordinates where z_g is the elevation of the sun, z_u is the elevation of the camera, θ_v is the declination of the camera from a normal vector to the surface, θ_0 is the declination of the sun from the same normal vector, ϕ_0 is the azimuth of the sun and ϕ_v is the azimuth of the camera.

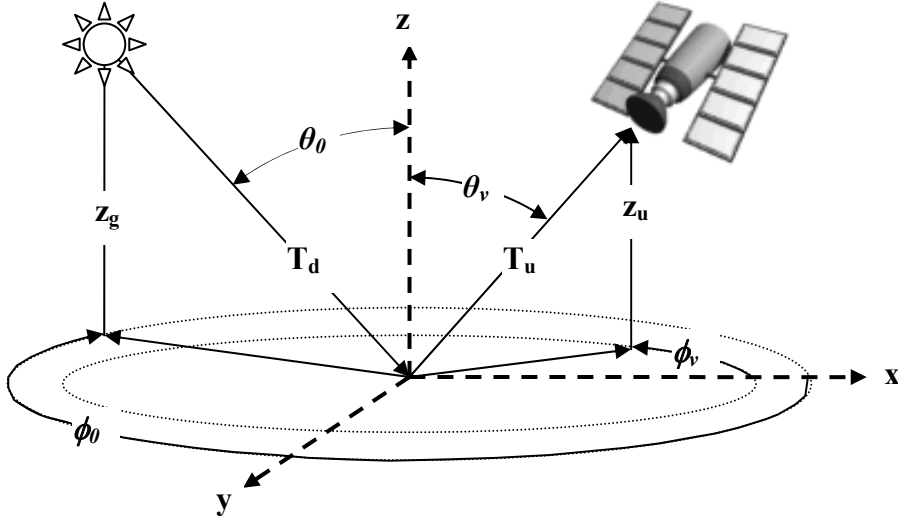


Figure 10: Source-Target-Receiver Geometry

3.1.1.1. Solar Spectral Signature E_0

For light to reach the sensor, light must first be generated. Ideally, the light source should be spectrally flat equally distributing the energy across all wavelengths. This can be accomplished in a laboratory setting, but in hyperspectral applications, the light source is typically the sun which has its own spectral signature. The sun's atmosphere is made of 73.46% hydrogen, 24.85% helium (by-product of the fusion of hydrogen atoms), and a fraction of other naturally occurring elements. These gases absorb certain wavelengths of light causing the documented Fraunhofer Absorption Lines [55]. Additionally, the fusion reaction produces more energy in the visible wavelengths. When these two effects are combined, it produces the typical solar spectrum seen in Figure 11. Thus, all images are colored with this solar spectrum.

The amount of sun light that reaches an object is a function of the sun declination angle and the downward atmospheric transmittance. The declination angle determines how much sun light directly hits an object. For example, when the sun is

directly overhead, the declination angle is zero and all the sun light reaches the object ($\cos(0^\circ) = 1$). When the declination angle is 60° , the amount of energy is only half of the energy when the sun is directly overhead. The interesting result of this effect is that the declination angle can be caused by either the sun being lower in the sky or the object sitting on a non-level surface. Thus, besides the angle of the sun relative to the horizon, even minor changes in topography can change the overall amount of sun light an object receives.

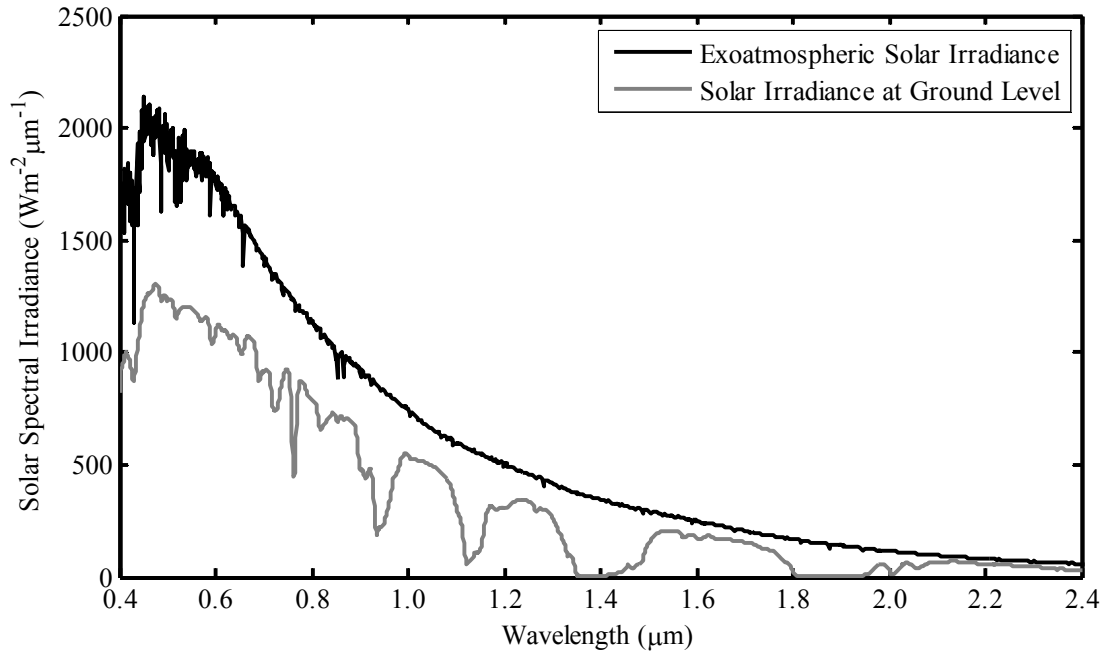


Figure 11: The Solar Spectrum

3.1.1.2. Downwelled Atmospheric Transmittance T_d

The other effect that reduces the sun light reaching an object is the downwelled atmospheric transmittance. The downwelled atmospheric transmittance quantifies the scattering and absorption effects that occur as light passes through the atmosphere. Scattering disperses the photons out of the direct path of the object thereby reducing the amount of light reaching the ground. The other dominant effect

is absorption which reduces the energy in certain wavelengths due to such molecules as water and carbon dioxide. By the time the light reaches the object being imaged, it has both the spectral properties of the sun and the intervening atmosphere as shown in Figure 11.

We can model how the atmosphere affects the sun light using a number of cylinders stacked on top of one another representing different altitudes. Each of these cylinders has a certain temperature, pressure, and humidity. These measurements dictate the amount of absorption and scattering that occurs within each cylinder and at each wavelength. Near the top of the atmosphere, there are very few particles and hence the three measurements are not as critical as near the bottom of the atmosphere. Thus, the cylinders are tall at the top of the atmosphere and become smaller as they reach the surface. This occurs because the dense atmosphere is located near the surface and causes a significant portion of the transmittance effects. This dense atmosphere is also the most variable as weather changes occur mostly in this region making signatures vary from one location to another.

3.1.1.3. Reflectance R

Once the sun light reaches the object, the reflectance of the object dictates which wavelengths of light are absorbed and which are reflected in various directions. The spatial reflectance attributes of a material are described by its bidirectional reflectance distribution function (BRDF). This function measures the reflectance for all wavelengths and input-output angles. A full BRDF characterization of a material is rare; so, materials are typically classified into gross categories ranging from specular reflectors to diffuse reflectors (also known as Lambertian). Specular materials reflect light in one direction such as mirrors. Diffuse reflectors reflect light

in all directions equally such as flat paint. Most materials fall between these two categories, but tend to be more diffuse than specular. Because BRDF characterizations are rare and most materials can be treated as diffuse, we assume diffuse reflectors for the remainder of this dissertation.

3.1.1.4. Upwelled Atmospheric Transmittance T_u

After the light has been reflected from the object being imaged, it passes back through the atmosphere to the sensor. The upwelled atmospheric transmittance quantifies these atmospheric effects. Upwelled atmospheric transmittance is very similar to downwelled atmospheric transmittance. The real difference between the two transmittances is upwelled transmittance only affects light between the object and the sensor. Therefore for low altitudes (e.g. 300m), this effect is minimized. On the other hand, the sensor could be space-borne in which case the light passes through the entire atmosphere. Either way, T_u is modeled the same way as T_d using cylinders of the atmosphere along the light path to quantify the scattering and absorption effects. As described in (1), the light reaches the sensor after being affected by the solar spectral signature, downwelled atmospheric effects, reflectance of the object being imaged, and upwelled atmospheric effects.

3.1.2. Sky Light

In the previous sections about atmospheric transmittance, scattering played an important part of how the spectral signature of the sun light was changed. This scattering of light has another side effect causing a secondary light source called sky light. Sky light can be mathematically modeled as

$$L_{sky}(x, y, \lambda) = R(x, y, \lambda) T_u(z_g, z_u, \theta_v, \phi_v, \lambda) \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi/2} E_s(\theta, \phi, \lambda) \cos \theta \sin \theta d\theta d\phi \quad (2)$$

where L_{sky} is the sky light radiance at the sensor, R is the reflectance of the object being imaged, T_u is the upwelled atmospheric transmittance, and E_s is the amount of energy scattered by the atmosphere.

Sky light takes a very similar path to sun light. Once the light reaches the object being imaged, it reflects the same as the sun light (assuming a diffuse material), and is reflected back up through the atmosphere to the sensor along the same path as the sun light. The main difference between sky light and sun light is the source of sky light is the scattering of photons in the atmosphere. These scattered photons arrive at the object being imaged from all directions. Therefore, these different patches of sky light are integrated over the hemisphere above the object being imaged. This produces the two integrals seen in (2) replacing the $T_d(z_g, \vartheta_0, \phi_0, \lambda)E_0(\lambda)\cos\vartheta_0$ term in (1).

There are three types of scattering that take place. The most well known scattering effect is Rayleigh scattering as explained by Lord Rayleigh to answer why the sky was blue [67]. Rayleigh scattering occurs when light interacts with the very small molecules that make up the atmosphere. The scattering occurs mostly in the blue wavelengths while other wavelengths are absorbed creating the blue color of the sky.

The other well known scattering effect is Mie Scattering [105]. This type of scattering occurs when photons interact with particles that are roughly the same size. These particles are typically composed of aerosols, combustible by-products, and small dust particles. This effect causes the scattered light around cities to be much different from the light scattered in rural areas.

The final effect is called non-selective scattering. This type of scattering occurs when the particles are much larger than the photons of light. Examples of such particles are water droplets and ice crystals that are due to cloud formations. Thus, scattered light can be affected by the amount and types of cloud cover in the image. These different scattering effects explain why images taken of rural areas on cloudless days can be very different from images taken of cities on partially cloudy days.

3.1.3. Upwelled Radiance

While some light is scattered so that it illuminates the object, other light is scattered directly towards the sensor. Unlike all the previous sources of light, upwelled radiance, L_{up} , never reaches the object being imaged. This light is scattered directly into the sensor's optical path from the atmosphere. Like sky light, it undergoes the same three scattering processes making it vary based on location and weather conditions. This has two effects on the imagery. The first effect reduces the overall contrast of the image. The second effect causes a blue shift (an increase in energy at the blue wavelengths) as the upwelled radiance term is typically dominated by Rayleigh scattering.

A good example of upwelled radiance is fog. As fog settles in, our eyes cannot see objects far away because they are obscured by the scattering of light towards our eyes from the water vapor particles (Mie and non-selective scattering). The effect is those objects disappear in a haze of gray. This effect is always present except it typically scatters such a small amount of photons relative to sun and sky light to make it undetectable in most situations.

The same can be said about the upwelled radiance reaching a sensor. In normal environmental conditions, upwelled radiance has a very small effect relative to the other sources of light. However, as the sensor is placed higher in altitude, the scattering effect becomes more predominant and can start to reduce the contrast of the image at the sensor. This occurs because there are more particles and thus more opportunities for scattering to occur.

3.1.4. Atmospheric Transfer Function

We can now mathematically define the radiance L reaching a sensor from an object with reflectance R as

$$\begin{aligned}
L(x, y, \lambda) = & R(x, y, \lambda)T_u(z_g, z_u, \vartheta_v, \phi_v, \lambda)T_d(z_g, \vartheta_0, \phi_0, \lambda)KE_0(\lambda)\cos \vartheta_0 \\
& + R(x, y, \lambda)T_u(z_g, z_u, \vartheta_v, \phi_v, \lambda) \int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi/2} E_s(\theta, \phi, \lambda) \cos \vartheta \sin \theta d\theta d\phi \quad (3) \\
& + L_{up}(z_g, z_u, \vartheta_v, \phi_v, \lambda).
\end{aligned}$$

The radiance equation in (3) states that the radiance at the sensor is a linear combination of the sun light, sky light, and upwelled radiance contributions. Although the final equation is a linear combination, the previous sub-sections detail how complex the atmospheric transfer function is to compute. Detailed weather information, source-target-receiver geometries, topography, and BDRFs are required to solve all the necessary functions. Typically, all of this information is not available and algorithms have to make simplifying assumptions. What assumptions are made depends on the type of algorithm.

3.2. Current Target Characterization Algorithms

Nearly all algorithms that convert reflectance to radiance or vice-versa are based on (3). The difference between these algorithms is the simplifying assumptions

they make and how they estimate each of the light sources. These algorithms can be broken down into two general methods: model-based methods and in-scene based methods.

3.2.1. Model-Based Methods

Model-based methods attempt to solve (3) directly. This type of solution requires a wealth of ancillary information besides the image. From Figure 10, the exact locations of the source, target, and receiver are required. This information is easy to obtain from the Global Positioning System (GPS). The location of the sun relative to a ground location is also well understood and can easily be found on the internet for a given location and time.

The information that is not as easy to obtain is weather data. In the modeling of atmospheric transmittance, the temperature, humidity, and pressure at varying levels of altitude need to be measured (i.e, the cylinders of the atmosphere). Typically, this is done using radiosondes. Radiosondes are weather sensors attached to balloons that measure all the needed weather information. Unfortunately, radiosonde information is not always available or applicable. For example, radiosondes are collected at certain locations which may be too far from the area being imaged to be applicable. If radiosonde data is available, the information is typically collected only twice a day and may describe the atmospheric profile that occurred hours in the past.

Murphy and Kolodner developed another way to get the requisite weather data [75]. If radiosonde data is not present or is inaccurate due to the aforementioned issues, weather maps generated from weather stations can be used. These weather maps produce an atmospheric profile that can be estimated via interpolation between

weather stations. This information is fused with satellite imagery to produce an accurate atmospheric profile at any location on the planet. This information is then used as the model inputs.

Once the ancillary information has been collected, a computational model can calculate the radiance for a given reflectance at any angle, source-target-receiver geometry, and wavelength via (3). MODTRAN is arguably the most used computational model [3]. It produces an estimate for every function in (3) and can make estimates for large declination angles as well as areas with variable topography. For most of the functions, it performs a direct calculation, but for the atmospheric transmittance functions, it has to make a simplifying assumption.

The scattering and absorption is not only a function of humidity, temperature, and pressure, but also of the constituent particles in the atmosphere. To model these particles in the atmosphere, MODTRAN uses one of many atmospheric profiles for urban, desert, or rain forest areas to name a few. Each profile uses a lookup table to provide an estimate of how light is scattered based on the types of particles found above each area type. Unfortunately, real world situations can vary significantly from the atmospheric profiles included with MODTRAN. While this may not greatly effect the radiance estimate, such assumptions can be very important when estimating weak target signatures such as Target 4.

Model-based methods have become the standard for atmospheric compensation techniques. They can make estimates for every parameter and function in the atmospheric transfer function. These estimates can take into account any type of topography and source-target-receiver geometry – even when the sensor may be on

or near the ground. To accomplish this calculation, they require a significant amount of ancillary information about source-target-receiver geometry, weather, and atmospheric profile type.

3.2.2. In-Scene Methods

The problem with model-based methods is that we sometimes lack all of the necessary ancillary information (or any estimate thereof). This is especially true with images collected in the past where such information was simply not collected. Because the information is either inaccurate or not available, another way to estimate the atmospheric transfer function was created using only the image data. These methods are called in-scene methods.

In-scene methods have to make a number of simplifying assumptions as well. The first assumption is that the area being imaged is small enough that the atmospheric profile (azimuths, altitudes, declination angles, etc.) is the same for all pixels even though this may not be true in a number of cases (e.g. water vapor [32]). The second assumption is that the pixels being used to estimate the atmospheric transfer function have Lambertian scattering properties. This assumption again is not necessarily true [89], but materials can be found that have near Lambertian properties that are acceptable for in-scene methods. Third, pixels that contain only one material (pure pixels) must exist in the image. Thus, in-scene methods are best for aerial images that cover a small amount of ground area.

3.2.2.1. Piech and Walker Shadow Method

One of the earliest and most accurate in-scene methods was developed by Piech and Walker [80]. They noted that shadow regions could be used to estimate the three main light sources in the atmospheric transfer function. Instead of estimating

detailed functions such as atmospheric transmittance, the atmospheric transfer function was simplified to

$$L(\lambda) = R(\lambda)L_{sun}(\lambda) + R(\lambda)FL_{sky}(\lambda) + L_{up}(\lambda) \quad (4)$$

where F is the fraction of the sky above the area being imaged (i.e., in shadow zones the amount of sky not blocked by the object creating the shadow). All x,y coordinates have been removed since we assume Lambertian scattering with equal amounts of light at each pixel.

The key to this method is realizing that in shadow zones, (4) becomes

$$L_{shade}(\lambda) = R(\lambda)FL_{sky}(\lambda) + L_{up}(\lambda) \quad (5)$$

since the sun light term has been reduced to zero. The algorithm therefore requires a material that is in both direct sun and shade conditions. When this occurs, the sunlight term can be easily calculated by taking the difference between (4) and (5) and solving for the sun light term to obtain

$$L_{sun}(\lambda) = \frac{L(\lambda) - L_{shade}(\lambda)}{R(\lambda)}. \quad (6)$$

To isolate the upwelled radiance term, equations (4) through (6) can be combined so the total radiance term is a linear regression of the shade radiance term as

$$\begin{aligned} L(\lambda) &= \frac{L_{sun}(\lambda) + FL_{sky}(\lambda)}{FL_{sky}(\lambda)} L_{shade}(\lambda) - \frac{L_{sun}(\lambda) + FL_{sky}(\lambda)}{FL_{sky}(\lambda)} L_{up}(\lambda) + L_{up}(\lambda) \\ &= m(\lambda)L_{shade}(\lambda) + b(\lambda). \end{aligned} \quad (7)$$

Using multiple materials with varying reflectance signatures, (7) can be solved to obtain the m and b terms at each wavelength. Rearranging these terms provides the upwelled radiance estimate

$$L_{up}(\lambda) = \frac{b(\lambda)}{1 - m(\lambda)}. \quad (8)$$

Equations (6) and (8) provide a way to establish the last light source such that

$$L_{sky}(\lambda) = \frac{L(\lambda) - R(\lambda)L_{sun}(\lambda) - L_{up}(\lambda)}{R(\lambda)F}. \quad (9)$$

This algorithm provides estimates of each light source within the atmospheric transfer function. The algorithm requires a shadow area which contains numerous pixels of the same material in both full sun and full shade conditions. Additionally, the algorithm requires multiple materials to be identified (historically by hand) to make estimates of the upwelled radiance term. In cases where these constraints cannot be met, we must rely on other methods.

3.2.2.2. Empirical Line Method

The empirical line method (ELM) is simpler than the shadow method and does not require any shadows in the imagery. ELM also does not estimate all of the light sources in the atmospheric transfer function. Instead, ELM makes the following simplification

$$L(\lambda) = R(\lambda)L_{sun+sky}(\lambda) + L_{up}(\lambda) \quad (10)$$

where the $L_{sun+sky}$ term combines the sun light and sky light into a single term assuming $F = 1$ due to the lack of shadows. Equation (10) identifies that the total radiance term is a linear combination of the upwelled radiance, the combined sun and sky light terms, and the reflectance. Thus, a linear relationship could be established by identifying a material with known reflectance in the scene. From this knowledge, the combined sun and sky light and upwelled radiance terms could be calculated for

each wavelength via linear regression. The linear regression is performed to estimate reflectance signatures from the radiance measurements in the image.

Various papers have identified numerous ways ELM can be implemented. All perform linear regression, but vary the number of materials required to estimate the parameters. The simplest implementations use one material and assume zero reflectance objects have zero radiance [26],[73]. This, of course, is not true as it assumes the upwelled radiance term simply does not exist. Not surprisingly, studies show errors of up to 20% in the predicted reflectance when compared to the true reflectance signature. Further studies used multiple known materials [26],[83] which show that four materials make the best estimates varying only a few percent from the actual reflectance signature.

While ELM has removed the need to have shadows, it does still require a significant number of known materials exist in the image. In cases where the study area is well documented or panels of known reflectance are placed in the scene, ELM performs very well. However, in images where only one material is well known, another method called dark object subtraction may be more applicable.

3.2.2.3. Dark Object Subtraction

Dark object subtraction is very simple. The idea is to find the minimum radiance values for each band in the image. These minimum values should represent the upwelled radiance assuming that the dark pixels have near-zero reflectivity. Using this dark object estimate as the upwelled radiance term allows the linear regression in ELM to take place without needing more than one known material.

This assumption holds in the NIR and SWIR bands, but the visible bands can have significant errors. The errors are especially troublesome when working with

subpixel targets which have low reflectance signatures. These low reflectance values from the targets inadvertently become part of the estimated upwelled radiance estimate. The overall effect in such cases is a corruption of the atmospheric transfer function and thus it is not well suited for subpixel detection.

3.3. Average Relative Radiance Transform

Another way to estimate the atmospheric transfer function is to use detection theory. There are a few reasons for approaching target characterization in this manner. First, the imagery does not have all the necessary ancillary information required by model-based methods. Second, the in-scene methods require user interaction to identify the materials with known reflectance in the image. This can be a time consuming process requiring a person with significant knowledge of remote sensing. Third, the simpler in-scene methods requiring the least amount of information are the most variable making them inappropriate for subpixel detection. Fourth, both in-scene and model-based methods were developed for analysis purposes. The idea was to map the radiances measured in the image back to reflectance values for comparison against spectral libraries for environmental research such as land class mapping and deforestation studies.

These reasons led us to develop a new atmospheric compensation algorithm for subpixel detection applications. To make subpixel detection applications accessible to a wide variety of users, the target characterization algorithm should automatically generate a target signature that can be used by a detector with little or no user intervention. The method should also use as little ancillary information as possible because this data may not always be available (e.g. historical image collections or analysis of areas for which information is not available). Finally, the

target characterization algorithm needs to provide enough fidelity that a detector can identify the target even among materials with similar spectral signatures.

The aforementioned constraints led us to develop the Average Relative Radiance Transform (ARRT). ARRT has a number of advantages. First, the algorithm is computationally efficient. Instead of projecting thousands (possibly millions) of pixels from radiance to reflectance, ARRT projects a few target reflectance signatures to radiance – a thousand or more so improvement in processing time. Second, ARRT is an in-scene atmospheric compensation technique requiring very little ancillary information. The algorithm only requires the image, the desired target reflectance signature, and a reference background reflectance signature. Source-target-receiver geometries and detailed weather information are not required. Third, ARRT is fully automated requiring only the aforementioned input signatures and image. Fourth, since ARRT is an in-scene method, the sensor need not be calibrated. As long as the errors in the sensor are uniform across the image, ARRT will account for the calibration errors where model-based methods cannot.

The original ARRT idea is based on the Internal Average Relative Reflectance algorithm (IARR) [59]. The IARR algorithm uses the spectral mean of an image as the atmospheric transfer function (ignoring upwelled radiance effects). The fundamental idea assumes that the image is comprised of many different underlying reflectance signatures that cancel one another when averaged together. The end result is the average spectral signature has a flat reflectance with some unknown multiplying factor K . Our early work demonstrated that applying IARR to generate target radiance signatures could work for subpixel detection algorithms [15]. The

drawback of the method is the assumption that the reflectance signatures cancel one another. Typically, the spectral mean still contains some of the reflectance characteristics of the dominant material. For example, if vegetation dominates the image, the spectral mean will have characteristics of the vegetation making it ineffective for certain targets.

This drawback led us to an updated ARRT algorithm that uses a two-pass detection method. The first detection pass identifies pixels with radiance values that most likely contain flat reflectances. This is very much like the underlying idea in IARR; however, ARRT directly detects these radiance signatures in the image instead of relying on the spectral mean.

To detect these highly probable flat reflectance materials in the image, a band ratio technique is employed. Band ratio techniques have been used in other analyses to identify vegetation, soil types, and other materials [48],[88]. For this application, we use a ratio between bands located on either side of the red-edge wavelength (700 nm). The red-edge effect causes a significant increase in reflectivity near 700 nm that corresponds to chlorophyll content (Figure 1) [90]. For radiance signatures generated from flat reflectance materials, the radiance drops slightly from 550 nm to 730 nm causing a band ratio less than one. Empirically, we found the value 0.8 to work best at identifying flat radiance signatures using both real-world HSI data and flat reflectance signatures generated by MODTRAN. Using this band ratio, radiance signatures with highly probable flat reflectances are found in the image and averaged together. As with IARR, the average reduces material and sensor variability to provide a better estimate of the flat reflectance than any single pixel found in the image.

To demonstrate the band ratio technique, we use the AVIRIS image in Chapter 2. For this data, we have two images with one being the true radiance measurement at the AVIRIS sensor and the other image being the estimated reflectance signatures for each pixel. The reflectance signatures were generated using model-based atmospheric compensation techniques validated by ground measurements of the scene [21]. Therefore, we will assume the reflectance estimates are accurate.

Figure 12 shows the results of the first stage of the ARRT algorithm on the AVIRIS data. In the top sub-figure, the mean spectrum of the radiance signatures chosen by ARRT to have highly likely flat reflectances is plotted. Using those pixel locations, we calculate the mean reflectance signature from the AVIRIS data in the second sub-figure. The reflectance is nearly flat across the spectrum except for some slight nonlinear effects near the lowest wavelengths. This slight decrease in reflectance is most likely an artifact of the AVIRIS reflectance estimation model. For example throughout the entire AVIRIS image, no one signature has a flat reflectance despite the presence of concrete in the image – a material with a known flat reflectance. Nevertheless, ARRT is finding radiance signatures that have a nearly flat reflectance signature.

The result of the first detection pass determines the spectral shape, but not amplitude. The average flat radiance signature is mathematically expressed as

$$L_{flat}(\lambda) = RL_{sun+sky}(\lambda) + L_{up}(\lambda) \quad (11)$$

where L_{flat} is the flat radiance signature estimated from the image. Because we assume the reflectance is flat, the reflectance term R should be constant for all

wavelengths. Additionally, L_{flat} includes the upwelled radiance term which causes a blue shift and loss of contrast as detailed in Section 3.1.3. Nevertheless, the L_{flat} term contains most of the spectral shape characteristics. Therefore, multiplying a reflectance signature by L_{flat} obtains a good representation of the spectral shape of the target material; however, the amplitude is still unknown as we do not have an estimate for R .

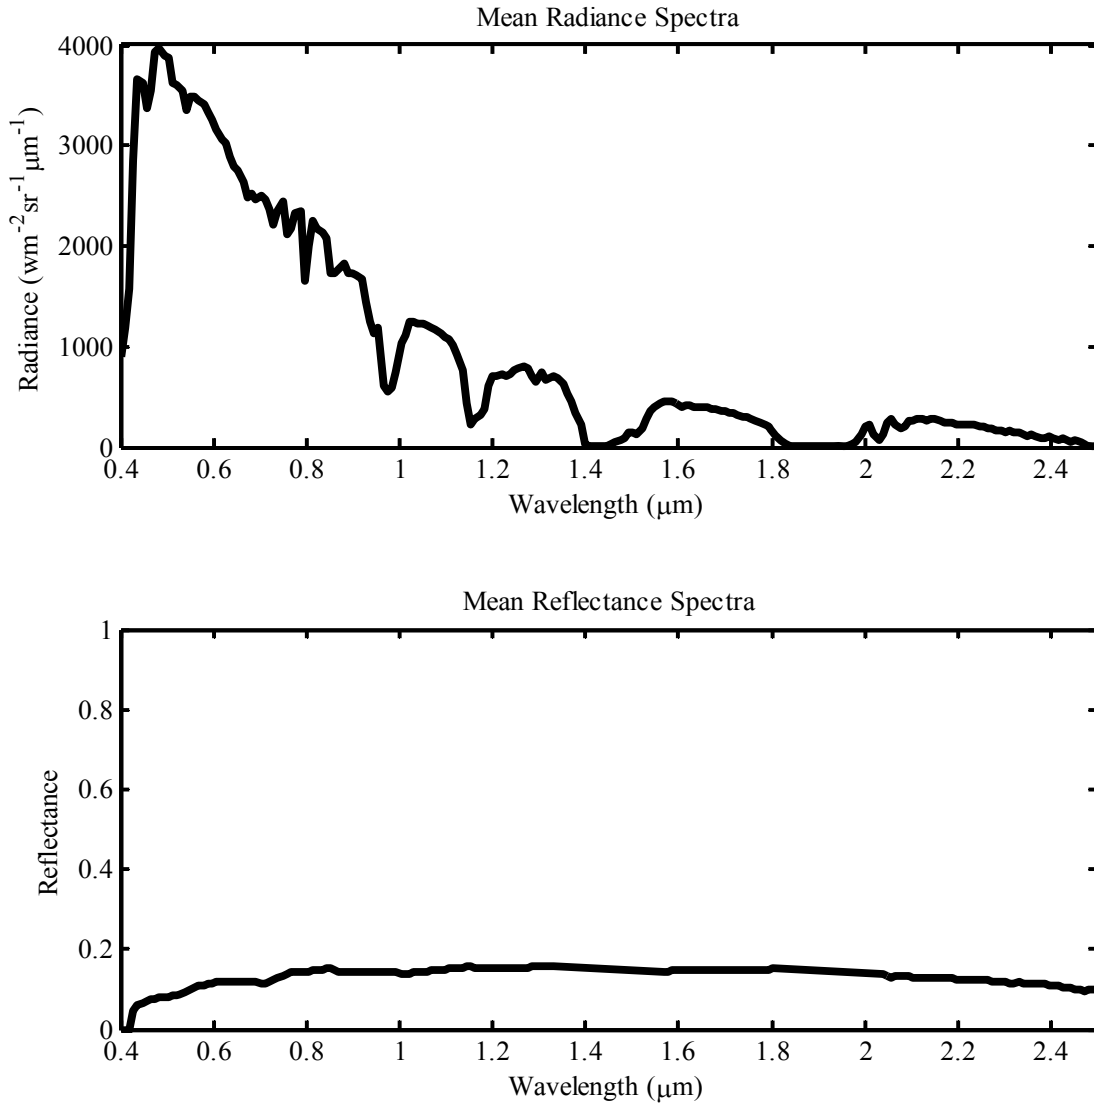


Figure 12: Comparison of Mean Radiance and Reflectance Estimates Using ARRT

It has been proposed that the amplitude mismatch is not problematic for detection applications. This statement is true in full pixel detection algorithms which use a replacement model (i.e., the pixel is either background or target, but not both). For full pixel target detection, the detectors normalize the pixels and desired target signature by their L^2 norm (see Spectral Angle Mapper [54],[95]). The result of such a normalization procedure makes the shape of the spectral signature the important determining factor as opposed to the amplitude. For replacement models, this is a desired result.

In subpixel target detection, the model is additive (i.e., the pixel is background or background plus target). To understand what happens if we divide a pixel by its L^2 norm, we describe a pixel using the linear mixing model introduced in Chapter 1:

$$x_{norm} = \frac{x}{\|x\|_2} = \frac{\sum_{i=1}^M \alpha_i e_i}{\left\| \sum_{i=1}^M \alpha_i e_i \right\|_2} = \frac{\sum_{i=1}^M \alpha_i e_i}{\left(\sum_{j=1}^M \sum_{i=1}^M \alpha_j \alpha_i e_j^T e_i \right)^{1/2}} \neq \frac{\sum_{i=1}^M \alpha_i e_i}{\sum_{i=1}^M \alpha_i \|e_i\|_2}. \quad (12)$$

Unlike full pixel targets, subpixel targets contain a number of background endmembers that are not a simple linear combination of their norms (i.e., cross terms exist in the solution). Therefore, normalizing the pixel, the background endmembers, and target spectra independently does not achieve the same result as full pixel target detection.

Because of this result, subpixel target detection requires a signature that is correct both in shape and amplitude. To estimate the amplitude, a second detection pass is required with a known reference material. Known reference materials refer to signatures within the image for which their reflectance signature is known. For

example, some ELM implementations use vegetation as a reference material. ARRT has no restriction on the reference material except that it has a moderate to strong reflectance signature and occurs as a pure pixel in the image.

A number of methods exist to choose a proper reference material. For example, reference signatures can be found based on the geographic region where the image was collected. If the image was collected over a desert region, sand would be an excellent reference signature while in forests, certain deciduous tree varieties would be a better match. All of these signatures are freely available from the United States Geological Survey (USGS) website (<http://speclab.cr.usgs.gov/>). Additionally, the USGS and other organizations have land class databases that describe the natural attributes of any area on the planet. From these two sources, a reference material for any image can be found.

Once a reference material and its corresponding spectral signature have been identified for the image of interest, the ARRT algorithm uses the Spectral Angle Mapper (SAM) algorithm to find the corresponding reference radiance signatures in the image [54]. Those pixels that pass a detection threshold are then ranked by their detection score. The top N detection scores are averaged to obtain the corresponding reference radiance signature for the image. Note we do not use the top N detection scores directly; instead we use the top N detection scores above a detection threshold. The reasoning behind this decision is that a given reference signature may not actually be within the image and the algorithm should not blindly use detection scores that fail to pass a minimum threshold. If there are no detections found in the image,

ARRT will inform the user and ask for another reference signature that better matches what is available in the image.

If a reference radiance signature is found, it is used to calculate the unknown reflectance R value in (11). The solution is

$$\hat{R} = \frac{\|L_{flat} R_{ref}\|_2}{\|L_{ref}\|_2} \quad (13)$$

where R_{ref} is the reflectance signature of the reference material and L_{ref} is the radiance signature estimated from the image for the reference material. R can be estimated assuming the reference signature has a high reflectance signature thus minimizing the effect of the upwelled radiance term.

An estimate of the upwelled radiance term can also be calculated as

$$\hat{L}_{up}(\lambda) = \begin{cases} \hat{R}L_{ref}(\lambda) - L_{flat}(\lambda)R_{ref}(\lambda), & \lambda < 700nm \\ 0, & \lambda \geq 700nm \end{cases} \quad (14)$$

The estimated upwelled radiance term is the difference between the estimated radiance signature and the detected radiance signature of the reference material in the visible wavelengths. In the near infrared and short-wave infrared wavelengths, errors due to noise dominate the signature. In the visible wavelengths, the Rayleigh and Mie scattering effects dominate, being significantly stronger than the error terms; thus, we clip the estimated upwelled radiance to only affect the visible wavelengths.

The final estimated target radiance signature can be calculated as

$$T(\lambda) = \left(\frac{L_{flat}(\lambda)}{\hat{R}} - \hat{L}_{up}(\lambda) \right) R_T(\lambda) + \hat{L}_{up}(\lambda) \quad (15)$$

where R_T is the reflectance signature of the desired target. To help clarify the ARRT algorithm, Figure 13 provides a block diagram describing the two-pass detection process and what inputs are necessary at each stage to arrive at (15).

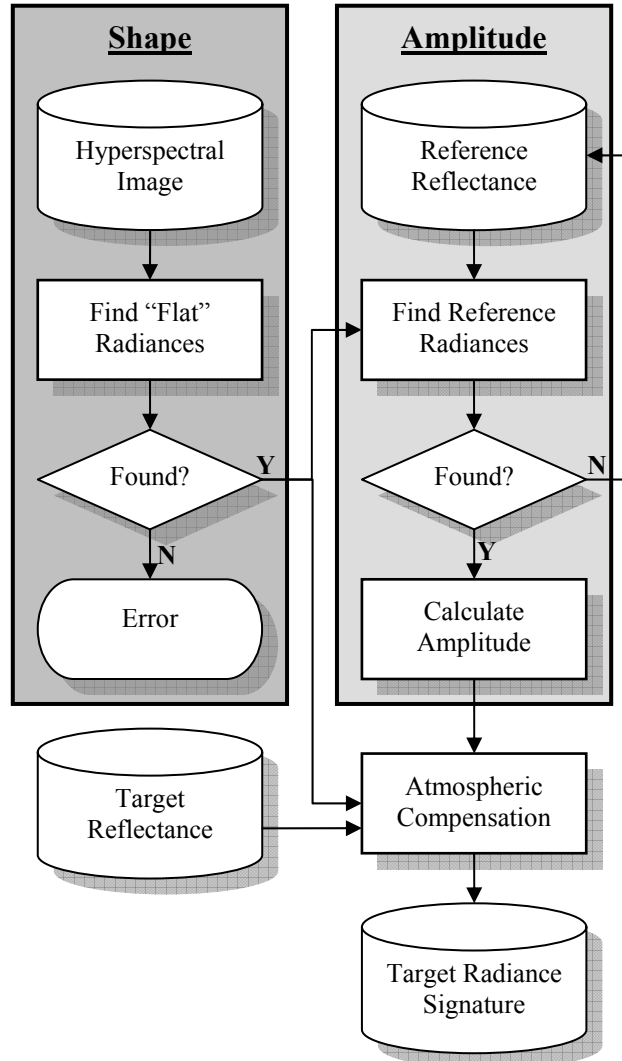


Figure 13: ARRT Block Diagram

Similar to other in-scene atmospheric compensation techniques, ARRT is only valid for certain conditions. First, ARRT was designed for aerial imagery where the upwelled radiance terms are small compared to the sun light and sky light terms. Second, ARRT requires a reference signature that has moderate to high reflectivity and has at least one pure pixel in the image. Currently, ARRT does not handle

shadow zones, but this can be addressed in another version that merges these techniques with Piech and Walker's work [80]. This will be discussed in more detail in Chapter 7.

3.4. Experimental Results

As with any atmospheric compensation algorithm, certain assumptions had to be made with ARRT. To validate whether these assumptions are valid and allow ARRT to produce useful target radiance signatures, we have designed two experiments. The first experiment uses Image 7 from Sensor X to directly compare target signatures generated by MODTRAN and ARRT to known target radiance signatures in the image. The second experiment compares target radiance signatures estimated using MODTRAN and ARRT relative to subpixel target detection performance.

Besides the imagery used for these experiments, a wealth of ancillary data was also collected. Radiosonde information was available from a nearby airport; however, this data was six hours old by the time the imagery was collected. Source-target-receiver geometry was also well documented as GPS was used on the airplane carrying the sensor. Numerous hand-held spectrometers were used on the ground to measure the reflectance of both target and background materials. While the sensor was not calibrated, the soil reflectance and radiance signatures were measured to correct for calibration errors via vicarious calibration as explained in Chapter 2. All of this ancillary data makes the following comparisons between MODTRAN and ARRT possible.

3.4.1. Comparison of Target Radiance Signatures

This experiment was used to validate the ARRT algorithm produces target signatures that match the actual target radiance signatures in an image. Image 7 from Sensor X was used for this experiment. The image was flown at 313m altitude so that each pixel imaged 0.0241 m^2 of area. The image contains Targets 3 and 4 with areas of 0.1090 m^2 and 0.0869 m^2 respectively. Targets thus spanned on average 4.5 and 3.6 pixels respectively.

Because the targets are multi-pixel, using the ground truth we received with the image, we were able to extract the true target radiance signatures from the image as shown in Figure 7 and Figure 8. These figures show the spectral variability of each target and their corresponding mean spectra. For Target 3, the mean spectrum is used in this experiment. For Target 4 however, we used only one signature pulled from a pixel that contained pure target spectra. Unfortunately, the smaller Target 4 only covers 3.6 pixels and thus has some background signature that “bleeds” into the target area as explained in Chapter 2. This minor corruption of the target signatures can be very serious when dealing with low reflectance targets. When the mean spectrum for Target 4 was used to test the subpixel detectors, it provided the worst detection performance supporting the hypothesis that many of the “true target” signatures were corrupted by background.

ARRT and MODTRAN were used to estimate Target 3 and Target 4 radiance signatures for Image 7. In the case of ARRT, two variants were used: one version estimated the upwelled radiance term while the other did not. The three estimated radiances were plotted against the known Target 3 and 4 radiance signatures in Figure 14 and Figure 15 respectively.

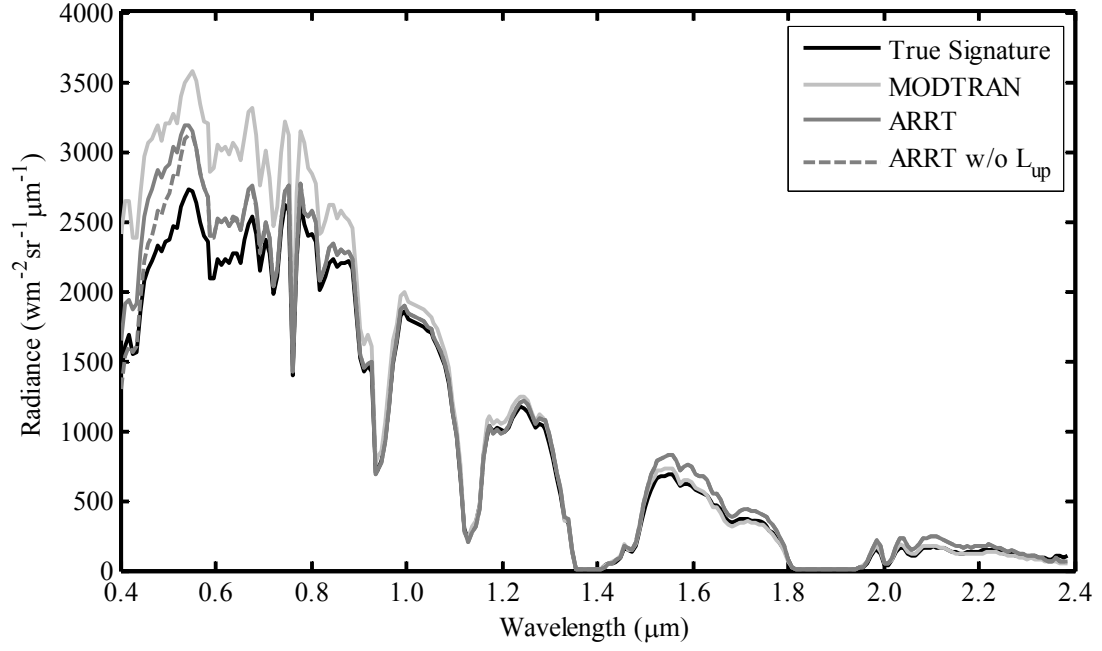


Figure 14: Comparison of Atmospheric Compensation Algorithms for Target 3

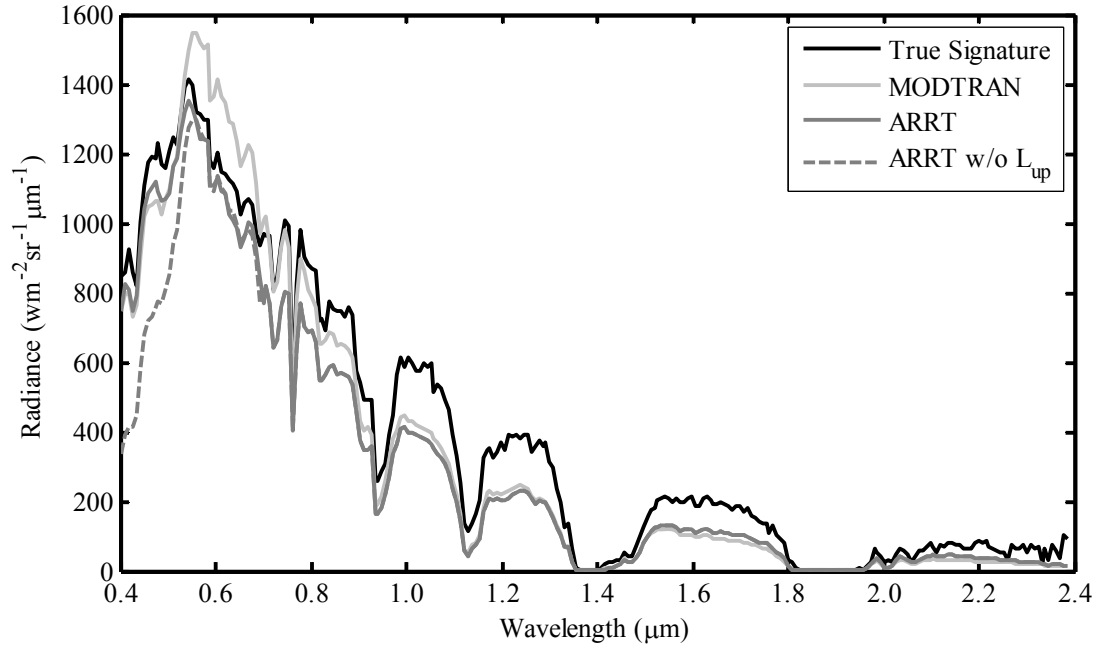


Figure 15: Comparison of Atmospheric Compensation Algorithms for Target 4

In addition, quantitative measurements are presented in Table 4. For each algorithm and target, two metrics were created measuring the similarity in amplitude

and similarity in shape to the true target signature. The metric for measuring the amplitude similarity is

$$\alpha = \|S - \hat{S}\|_2 \quad (16)$$

The metric for measuring the shape similarity is the angle between the spectral signatures

$$\theta = \cos^{-1} \left(\frac{S^T \hat{S}}{\|S\|_2 \|\hat{S}\|_2} \right). \quad (17)$$

The estimated target radiance signature that minimizes the above metrics provides a better match to the true target radiance signature.

Table 4: Quantitative Comparison of Atmospheric Compensation Algorithms

Target	Metric	MODTRAN	ARRT	ARRT (No L_{up})
3	α	5664	2547	1888
	θ	5.86°	4.00°	2.98°
4	α	1515	1648	2342
	θ	9.81°	7.83°	10.49°

Comparing the signatures using Figure 14, Figure 15, and Table 4, ARRT estimates the target radiance signatures well. For Target 3, ARRT outperforms MODTRAN in matching the true target signature. The shape and amplitude is a better match and as such we expect to have better detection performance using the ARRT signature. Interestingly, the ARRT version without an upwelled radiance is marginally better than the standard ARRT algorithm.

For Target 4, the results are mixed. MODTRAN estimates the amplitude very well, but does not do as well estimating the overall shape of the signature. The ARRT algorithm estimates the shape better than MODTRAN, but underestimates the

amplitude. The ARRT algorithm without the upwelled radiance term performs the worst of all the variants. All algorithms however underestimate the shape and amplitude of the SWIR bands including MODTRAN. In the next section we show that this underestimation will lead to poor detection performance. Thus, Target 4 is an interesting case for further research into ways to improve all atmospheric compensation techniques.

Overall the ARRT algorithm performs as well as MODTRAN using only the target reflectance signature, reference signature, and imagery. MODTRAN requires radiosonde information, vicarious calibration, and GPS information to produce signatures that are at best only slightly better than ARRT. Considering the amount of time necessary to collect all this information and process it through MODTRAN, ARRT provides similar target estimates with significantly less ancillary data and in a fraction of the time.

3.4.2. Comparison of Target Signatures for Subpixel Detection

While comparing the estimated radiance signatures to their true counterparts is important, it does not answer whether the estimated targets are a good match for subpixel target detection applications. This set of experiments was designed to answer the aforementioned question using the well known Adaptive Cosine Estimate (ACE) algorithm [58]. This detector is one of the better detectors available for subpixel detection in HSI data. Another reason for using this detector is the background is modeled entirely by a multivariate normal distribution; thus, no background endmembers are required. The algorithm's performance is based solely on the image and the target signature. Thus, ACE makes an ideal algorithm to use for experiments

comparing algorithms that generate target radiance signatures. More information on the ACE algorithm is documented in Chapter 5.

For all of these experiments, the ACE algorithm was processed in the following manner. Besides the target signature, a mean and covariance had to be estimated. There are two ways to estimate these parameters: globally or locally. We chose the global method for these experiments as this provided both the best performance and the fastest implementation. Typically, the SAM algorithm is used to detect obvious target detections and remove them from the image before calculating the global mean and covariance as was done for Image 7. In Images 1 through 6 however, the targets are so small, they are not detected by the SAM algorithm and hence were not removed. While this may slightly degrade performance [27], it provides the most honest performance results as real-world applications will not have knowledge of the ground truth a-priori.

Once the ACE detector was run, a detection image was generated. As mentioned in Chapter 2, the ground truth for Sensor X was for object level detection. To obtain objects from our detection images, a clustering threshold is applied. This clustering threshold refers to a threshold that combines adjacent pixels together to form an object which will be classified as either target or clutter. Typically this threshold is chosen to include no more than 1% to 5% of the pixels in the image depending on the application. In our analysis, we chose 1% as we knew the number of targets was far less than 1% of the pixel in any one image. Each cluster is assigned the maximum detection score from all the pixels that make up the cluster. Along with the maximum detection score, each cluster is identified as either target or clutter

based on their location relative to the object-level ground truth. This information can then be used to identify how well a detector performs.

3.4.2.1. Comparison of Full Pixel Detection Performance

The first experiment applies ACE to Targets 3 and 4 in Image 7 from Sensor X using the target signatures generated in the previous set of experiments. For this experiment we use the MODTRAN algorithm and three variants of ARRT: the standard ARRT algorithm described in the previous sections, the ARRT algorithm without the upwelled radiance estimate (ARRT w/o L_{up}), and an adjusted ARRT algorithm where the amplitude has been matched perfectly to the extracted target signatures (ARRT Adj). The ARRT variants were added to identify the benefits of estimating the upwelled radiance term and to test the importance of obtaining a correct estimate of amplitude.

Figure 16 shows the ACE detector results for the estimated target signatures. Each figure contains black and gray vertical bars. The black bars show the range of detection values for the background. The gray bars show the range of detection values for the targets. Ideally, these bars should not overlap indicating the targets are completely separable from the background. Above the black bar, a number is posted identifying how many false alarms occur above the minimum target detection score (i.e., the number of false alarms that are in or above the range of target detection scores). Above the gray bar, a number is posted indicating the percentage of target detected in the image.

Results for Target 3 show all the target estimates are well matched to the targets in the image. The ARRT estimates achieve the ideal case separating the target from the background easily. The MODTRAN signature generated 4 false alarms, but

this was to be expected as it was not as accurate in both shape and amplitude as the ARRT signatures. Even with 4 false alarms, the performance is only marginally worse than using the ARRT signatures.

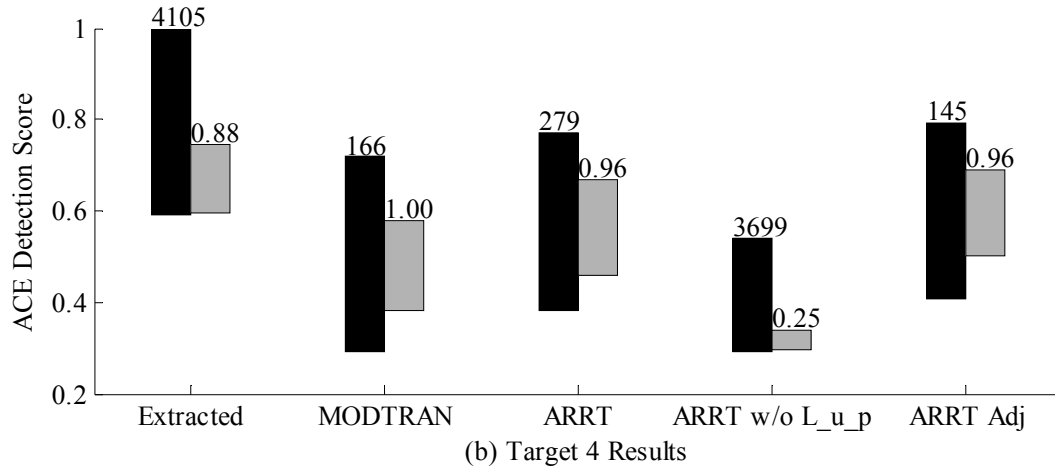
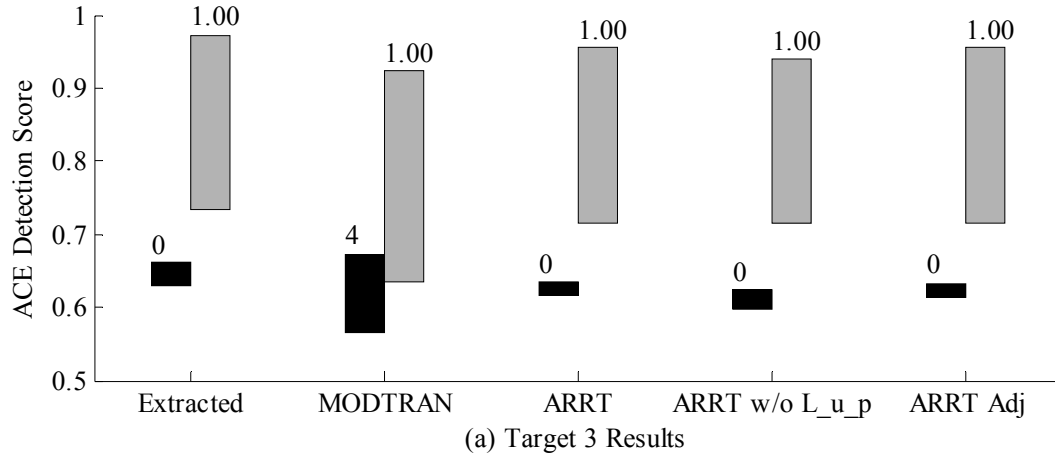


Figure 16: ACE Results for Image 7
for (a) Target 3 and (b) Target 4

Results for Target 4 are much more interesting. First, Target 4 is a difficult target to detect because of its low reflectance signature. Not surprisingly, the false alarm counts are significantly higher with this target than with Target 3. The MODTRAN signature provides the best performance outperforming the “true” signature estimated from the mean of the target detections in the image. ARRT provides good detection performance, but has 68% more false alarms. As expected,

the ARRT estimate without the upwelled radiance term performs the worst giving an abysmal 25% P_d providing evidence that the upwelled radiance term is important to subpixel detection applications. Another interesting result is the last set of bars. These results were generated using an ARRT signature that was corrected to have the same amplitude as the target signature taken from the image. The results for this signature rival the performance achieved with MODTRAN. Thus, amplitude plays a considerable role in achieving good subpixel detection performance.

On a final note, the true target signature for Target 4 does not perform as well as most of the target radiance estimates. This is not surprising however given the size of Target 4 in Image 7. Since targets span only 3.6 pixels, most likely some “target” pixels were identified that contained some background materials as well. Thus, the “real” target signature is compromised and this leads to the degraded performance. Another result from this experiment is that even with multi-pixel targets that contain few pixels; atmospheric compensation algorithms may provide a better estimate of the target than can be drawn from the image with known ground truth.

3.4.2.2. Comparison of Subpixel Detection Performance

Image 7 provided us the opportunity to compare target signatures generated using atmospheric compensation algorithms to their true signatures in an image. Unfortunately, the analysis could not provide performance estimates for actual subpixel targets. To provide this type of analysis, we compare the MODTRAN, ARRT, and ARRT without L_{up} on Images 1 through 6 from Sensor X. These images were collected at an altitude of 1220 m so that each pixel imaged approximately 0.1820 m^2 of area. The result of the higher altitude is that the targets have fill factors

(percent of the pixel occupied by target material) ranging from at most 60% to as low as 11%.

As was done in the previous experiment, ACE was applied to the data for the various target types and target radiance estimates. A clustering threshold of 1% was used to form the objects that were identified as either target or clutter using the provided ground truth. Some target did span multiple pixels, but did so with smaller fill factors (e.g., Target 3 has a 60% fill factor that can be split across two pixels as 20% and 40%).

Instead of bar graphs to analyze performance, receiver operating characteristic (ROC) curves were used. These ROC curves were generated across all images so enough targets would be available to make a meaningful ROC curve. As is typical, the y-axis measures the P_d normalized to 1. The x-axis, however, is a measure of false alarm density. This metric is the number of false alarms divided by the total area imaged. Curves for detectors that achieve false alarm densities of 10^{-3} or lower with 50% P_d are considered good performers.

Figure 17, Figure 18, and Figure 19 display the ROC curves for Targets 1 through 3 respectively. In all cases, ARRT performs as well as MODTRAN. This shows that an in-scene technique can perform as well as a complicated model-based technique for subpixel detection performance. This result is expected given the good results seen on Target 3 in the earlier experiments. Additionally, Targets 1 through 3 have moderate to strong reflective signatures as shown in Figure 6. Because the signatures have good reflectance, the algorithms are less prone to small errors and provide good radiance estimates in all cases.

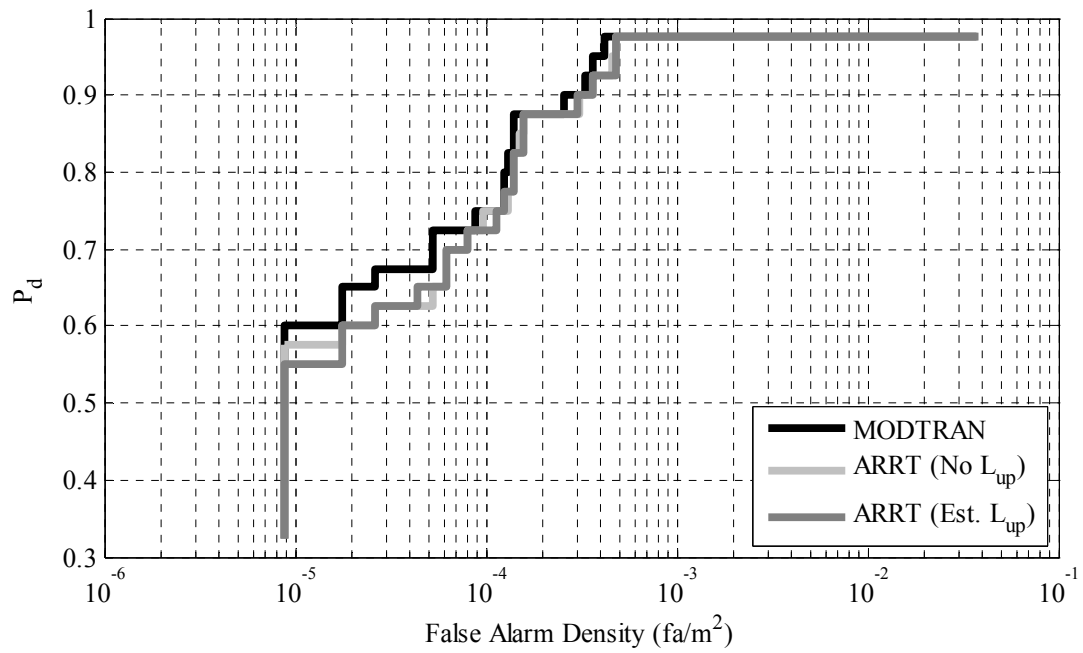


Figure 17: ROC Comparison of Target 1 Signatures

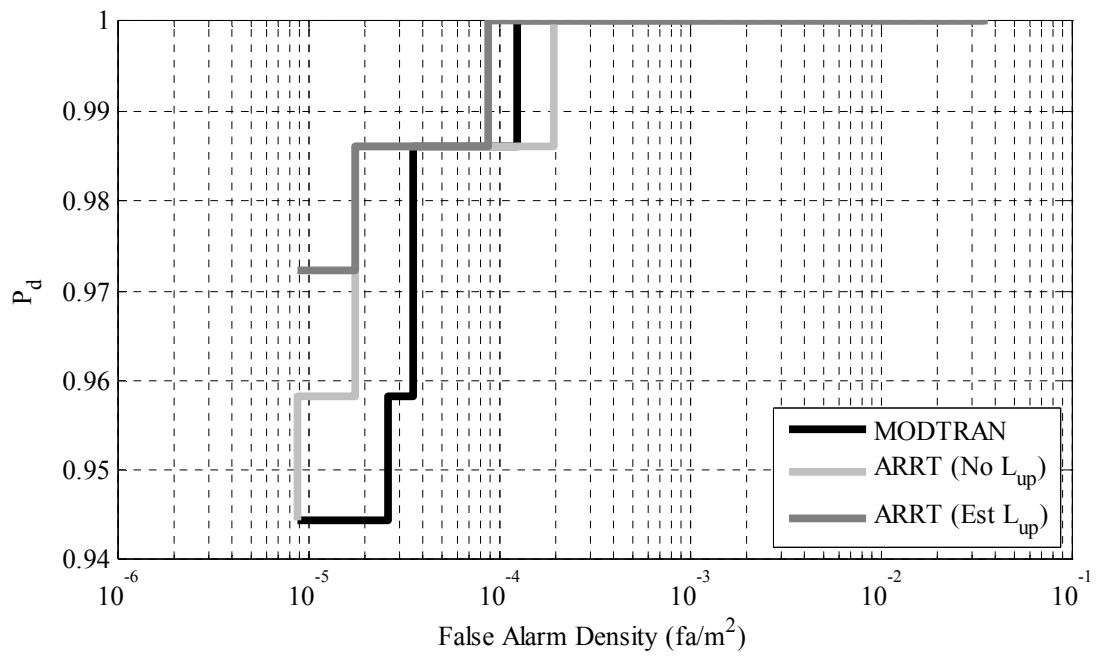


Figure 18: ROC Comparison of Target 2 Signatures

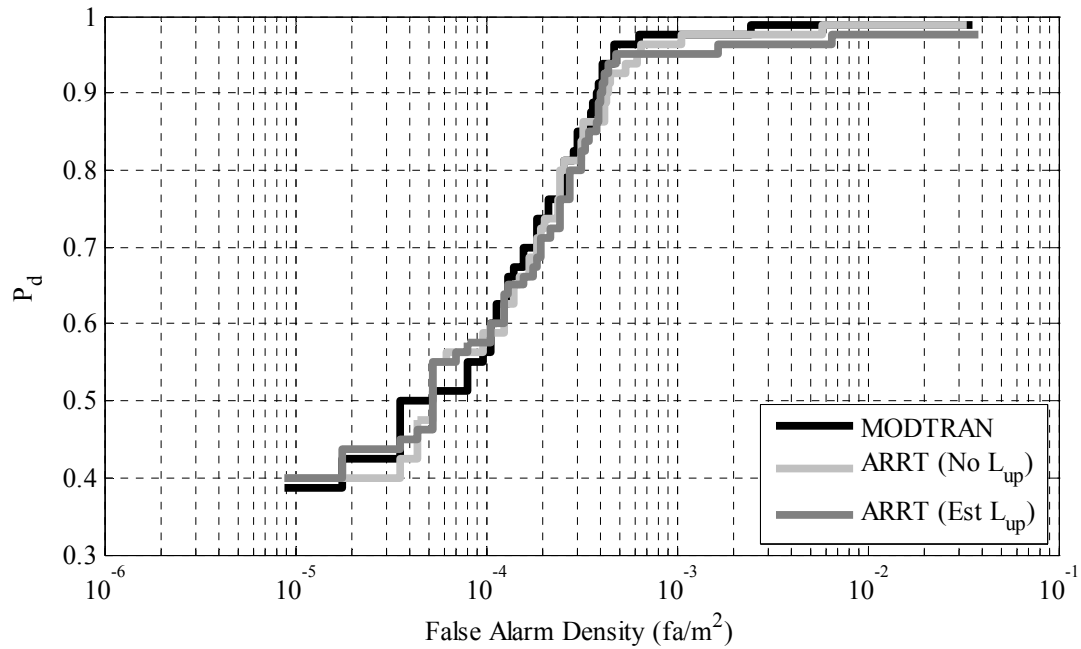


Figure 19: ROC Comparison of Target 3 Signatures

Target 4 is the difficult target. As mentioned in the previous section, the target has a weak reflectance signature making it hard to detect at an altitude of 313m. At 1220m altitudes, the target becomes very difficult to detect. None of the detectors with any target estimate perform well although MODTRAN performs the best as expected. Model-based methods are somewhat immune to sensor collection errors and tend to perform better with low reflectance targets [93]. In-scene methods tend to degrade with such targets as even small errors can seriously affect the shape and amplitude of the estimated target signature which leads to degraded detection performance. Therefore when dealing with weak target signatures, model-based methods still have an advantage over in-scene methods as has been previously documented [93]. This statement holds true for ARRT as well.

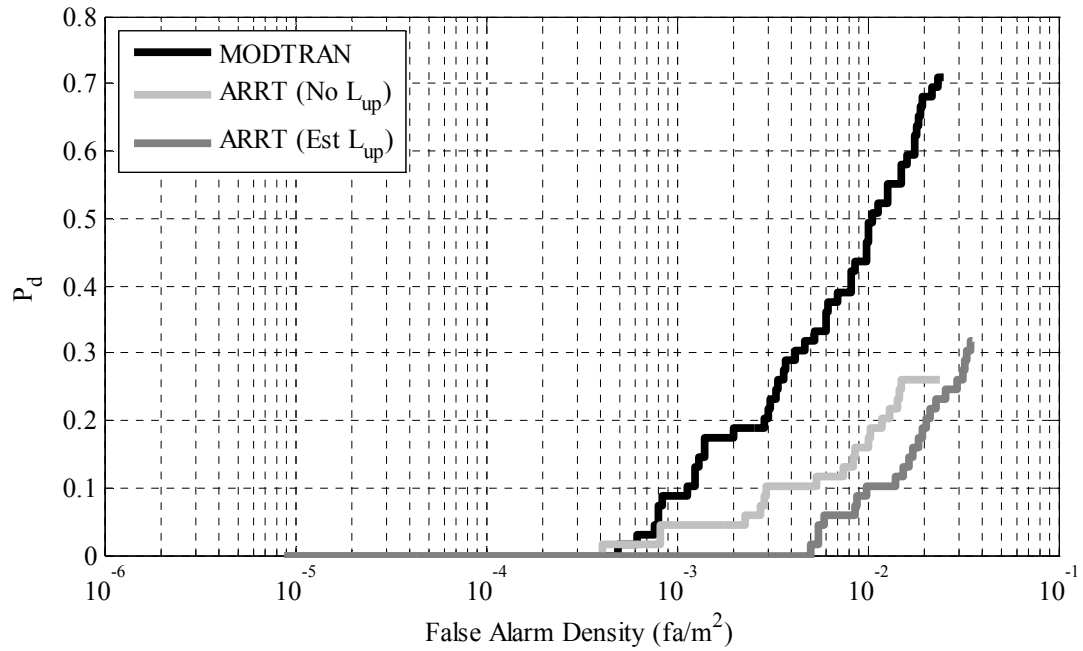


Figure 20: ROC Comparison of Target 4 Signatures

3.5. Summary

Characterization of the target radiance signature is a key part of subpixel detection. Many ways have been developed over the years to estimate the atmospheric transfer function at the heart of target characterization. This work presents a new in-scene algorithm ARRT for characterizing target radiance signatures using only the image and a reference reflectance signature. The algorithm uses detection theory and radiative transfer theory to project a target reflectance signature into the radiance seen at the sensor.

The ARRT algorithm provides a number of advantages over other methods. First, ARRT provides radiance signatures in a fraction of the time of model-based methods since ancillary information such as weather and source-target-receiver geometry are not used. Second, ARRT generates signatures that rival model-based methods. Third, the signatures generated by ARRT have been shown to provide good

subpixel detection performance over a variety of targets. Finally, sensor calibration issues which are problematic for model-based methods pose no problem for in-scene methods such as ARRT. These traits make ARRT very attractive for applications where a model is simply not feasible and or the ancillary information cannot be obtained.

While ARRT does have the aforementioned attractive properties, it also has its limitations. ARRT is meant for aerial imagery as opposed to satellite data or images taken at extreme oblique angles. Additionally, the imagery must contain pure background pixels with moderate to high reflectance signature to estimate the amplitude of the target radiance signature. As expected, ARRT like other in-scene methods has difficulty estimating signatures with low reflectance. However, even in this extreme case, model-based methods perform only marginally better.

Chapter 4: Background Signature Characterization

Target characterization is an important aspect of any detection algorithm. In subpixel detection, however, characterization of the competing background signatures within the pixel is just as important. Unlike conventional full-pixel detection where the pixel contains target or background signatures, subpixel targets are a combination of the target and the competing background signatures as described by the linear mixing model in Chapter 1. Having developed a way to characterize the target signature, we must now focus our attention on characterizing the background signatures.

Unlike target characterization where we have a known target signature, we do not know a-priori all of the background materials in an image. Instead, these background materials must be estimated. While one could use land class maps to identify the main background components in any area, these maps are typically coarse and cannot capture the material variability that may be in the scene. Thus, most subpixel detectors rely on endmember extraction methods which adaptively estimate the background endmembers from the image.

This chapter begins by providing an overview of endmember extraction techniques. The first section describes some of the many endmember extraction techniques available to the community today. While this is not an exhaustive list, it does provide examples of the fundamentally different ways background endmembers can be estimated. From this list, we identify the two endmember extraction techniques we use for the remainder of the dissertation and motivate why we selected them.

In the following sections of the chapter, we discuss the importance of estimating the correct number of endmembers for subpixel target detection purposes. We argue that this topic has been largely ignored by the community based on the different ways researchers have estimated the number of endmembers. We introduce the various state-of-the-art methods from intrinsic dimensionality to virtual dimensionality statistics. We present two of our own proposed methods for estimating the number of endmembers arguing that the estimate should be based on both the endmember extraction algorithm and the desired target signature. We compare our methods to the current state-of-the-art methods showing appreciable gains in a number of experiments. Through these comparisons, we also show how important correct estimation of the number of endmembers is to subpixel detection performance.

4.1. A Review of Endmember Extraction Methods

A number of algorithms have been developed to adaptively estimate the endmembers in an image. A review of the literature shows how many different algorithms exist including Pixel Purity Index (PPI) [9], N-FINDR [116], the Simulated Annealing Algorithm (SAA) [7], Optical Real-Time Spectral Identification System (ORASIS) [37], Iterative Error Analysis (IEA) [77], and Automated Morphological Endmember Extraction (AMEE) [81] to name just a few. A good review of various endmember extraction algorithms can be found in [82]. The intent of this section is to simply and quickly describe the different ways endmembers can be extracted from HSI data.

The PPI, N-FINDR, and SAA algorithms are geometry-based methods. These algorithms project the HSI data into a smaller dimension d using methods such as the

Maximum Noise Fraction (MNF) transform [36]. After the transformation, the algorithms have slightly different approaches. PPI generates random lines onto which the transformed data is projected. The outliers on each line are counted and the process is repeated many times identifying those pixels that continue to be outliers as endmembers. An operator takes this result and uses a d -dimensional visualization tool to identify the final number of endmembers. N-FINDR finds the endmembers as the $d+1$ vertices of the simplex that contains the maximum amount of the transformed data. N-FINDR is computationally efficient and can be performed in near real-time without operator intervention. SAA is very similar to N-FINDR in that it also identifies endmembers as the vertices of a simplex enclosing the transformed data. Unlike N-FINDR though, SAA creates “virtual endmembers” when no pure pixels are present in the image. This generation of virtual endmembers using a simulated annealing algorithm guarantees endmembers that are pure material spectra. This is also an automatic extraction technique, but is more computationally expensive than N-FINDR due to the simulated annealing.

ORASIS is both a vector quantization method and geometric method. This algorithm developed by the U.S. Naval Research Laboratory (NRL) operates in real-time using a two step process. The first pass reduces the volume of the HSI data using a learning vector quantization (LVQ) process [10]. Using LVQ, exemplar signatures are adaptively found from the image using a distance metric (typically the SAM metric [54]). Once the exemplars are found, a modified Gram-Schmidt process called salient selection is used to project the exemplars onto a smaller dimensional orthogonal subspace. The algorithm identifies the endmembers as those that make up

the vertices of the simplex that encloses the projected data similar to N-FINDR and SAA; however because of the LVQ preprocessing step, this algorithm can run in real-time.

AMEE is a joint spatial and spectral morphological approach to endmember extraction. In this method, no subspace projection is necessary. Instead, the image is iteratively processed using spatial morphological kernels of various sizes. At each pixel location, the spectrally purest and spectrally most mixed pixels are found. The morphological eccentricity index (MEI) is calculated as the angles between these pure and mixed pixels. This is repeated for multiple kernel sizes until an MEI image is created. Segmentation takes place on the MEI image and the endmembers are those chosen from the image after a spatial and spectral growing procedure occurs which removes variability within each spectral class.

The IEA algorithm extracts physically meaningful endmembers that are based on minimizing the mean squared error between the actual image and an unmixed image. The algorithm begins with the target signature and unmixes the image (estimates the endmembers and corresponding abundances) using the Fully Constrained Least Squares algorithm [46] (further details can be found in Chapter 5). An error image is created between the original image and the unmixed image generated using (1). The mean of the pixels that contain the largest mean squared error are chosen as the next endmember. Extraction continues until N number of endmembers is found.

There is another class of endmember extraction methods based on statistical models. Parametric statistical models include the stochastic mixing models (SMM)

[104] based on expectation maximization methods and the Modified Spectral Mixture Analysis (MSMA) which is an approach similar to the SAA algorithm [110]. Non-parametric statistical algorithms have also been used to extract endmembers such as K-Means clustering [29].

4.2. Selected Endmember Extraction Techniques

To characterize the background for subpixel target detection, we are interested in finding an endmember extraction technique that 1) performs well, 2) produces physically meaningful endmembers, and 3) is fully automatic. Using the research from [77] and [82], we decided on a variant of the IEA algorithm for multiple reasons. First, the IEA algorithm produces physically meaningful endmembers that are well matched to the FCLS algorithm – an abundance estimation algorithm that will be used in our subpixel detectors described in Chapter 5. Second, the algorithm provides endmembers that are significantly different from the target signature minimizing the change of background signatures “bleeding” into the target subspace. Third, the algorithm runs quickly taking only a few minutes to extract 30 endmembers. Fourth, the IEA algorithm was identified as one of the best performing endmember extraction techniques in [82]. Since the IEA algorithm is also fully automatic, it meets all of our criteria.

We use another technique defined by the popular Adaptive Matched Subspace Detector (AMSD) – a baseline subpixel detector used in Chapter 5. We use this method because the AMSD algorithm specifically identifies this method be applied in its detector [71][109]. This technique does not extract physical endmembers. Instead it performs an eigenvector decomposition of the image correlation matrix. The resulting eigenvectors comprise the endmembers for the background. Note that while

these endmembers are not physically meaningful, they do minimize the mean squared error when used with the AMSD algorithm. We only use this method for the AMSD algorithm as it does not provide physically meaningful endmembers for our physics-based approach.

4.3. Dimensionality of Hyperspectral Imagery

In addition to the extraction of endmembers, a significant amount of research has gone into identifying the correct number of endmembers for a scene. Most algorithms have focused on what has been termed “intrinsic” dimensionality [19]. These dimensionality measures focus on identifying the unique spectral signatures in an image. For classification purposes, it is important to estimate the intrinsic dimensionality. For target detection applications, intrinsic dimensionality may not be the best measure.

In target detection, the background must be characterized such that the probability of detecting the target is maximized while the probability of detecting a false alarm is minimized. In such cases, the number of endmembers required to characterize the background may be significantly more than the intrinsic dimensionality. The reasons are varied, but can be quickly summarized as the additional endmembers may be signatures due to shadowing effects, sensor artifacts, and finer material identification (e.g. coarse sand vs. fine sand). This has been noted in [19] where the best number of endmembers varied for different applications. This measure of dimensionality relative to detection performance has been termed virtual dimensionality [19].

The next two sections describe the different metrics used to select the “best” number of endmembers from a scene. The intrinsic dimensionality measures are

energy, Akaike Information Criterion (AIC), Minimum Description Length (MDL), and Empirical Indicator Function (EIF). The virtual dimensionality measures are based on work by Chang and Du [19], Thai and Healey [109], and two we propose for subpixel detection applications.

4.3.1. *Intrinsic Dimensionality Metrics*

4.3.1.1. *Energy Metric*

This metric is used by Manolakis, Siracusa, and Shaw for the AMSD algorithm [71]. In this paper, they characterize the background as the eigenvalue decomposition of the image correlation matrix. The resulting eigenvalues are sorted in decreasing order. The number of endmembers used is calculated using the sorted eigenvalues such that

$$\hat{m} = \min_m \left(\frac{\sum_{i=1}^m \lambda_i}{\sum_{j=1}^M \lambda_j} \geq 0.99 \right) \quad (18)$$

where M is the total number of endmember extracted and λ_i is the i^{th} ordered eigenvalue.

4.3.1.2. *MDL Metric*

A set of metrics was developed to estimate the order of a statistical model. One of the first was the AIC published by Akaike in 1974 [2]. The AIC statistic was found to be inconsistent [51] and this led to other works by Rissanen using an information-theoretic criterion [87] and by Kashyap [50] and Schwartz [96] using a Bayesian framework. The researchers independently came to the same result: the Minimum Description Length (or Bayesian Information Criterion as Schwartz identified it). The criterion is

$$\hat{m} = \min_m \left(-\log L(x, \alpha_m) + \frac{1}{2} k \log N \right) \quad (19)$$

where $L(x, \alpha_m)$ is the statistical likelihood function parameterized by α_m , k is the number of free parameters that must be estimated, N is the number of samples used to estimate the likelihood and its associated parameters, and m is the dimension of the parameters.

Chang and Du used Wax and Kailath's MDL criterion in their research [113]. The results showed poor performance because of two reasons. First, the Wax and Kailath work was designed for time series data where each sample came from an iid zero-mean Gaussian distribution; therefore, the combined likelihood could be expressed entirely in terms of the data covariance matrix. HSI data does not fit this assumption as mentioned in [71] and [103]. Second, Chang and Du used the equation directly from Wax and Kailath [113] which was designed for complex data. HSI data is real-valued and hence the equation they used was inappropriate. Instead, the equation should have been

$$\hat{m} = \min_m \left(-\log \left(\frac{\sum_{i=1}^L \lambda_i^{1/(L-m+1)}}{\frac{1}{L-m+1} \sum_{i=1}^L \lambda_i} \right)^{(L-m+1)N} + \frac{1}{4} m(1+2L-m) \log N \right) \quad (20)$$

where λ_i are the eigenvalues of the image covariance matrix, L is the number of spectral bands, and N is the number of pixels in the image. Nevertheless, in all of our experiments, the Wax/Kailath implementation never achieved a minimum (using Wax and Kailath's original equation or (20)).

4.3.1.3. EIF Metric

Malinowski created a metric specifically designed to estimate the number of unique spectra in chemical spectroscopy studies [68]. Using empirical studies based on chemical factor analysis, he created an empirical indicator function (EIF) such that

$$\hat{m} = \min_m \left(\frac{\left(\sum_{i=1}^M \lambda_i \right)^{1/2}}{N^{1/2} (L - m + 1)^{3/2}} \right) \quad (21)$$

where λ_i are the eigenvalues of the $L \times M$ endmember matrix, M is the total number of endmembers, L is the number of spectral bands, and N is the number of pixels in the image.

4.3.2. Virtual Dimensionality Metrics

4.3.2.1. NSP Metric

The term “virtual dimensionality” was coined by Chang and Du [19]. In this paper, they presented a new way to assess the dimensionality of HSI data relative to classification and detection performance. Interestingly, the Noise Subspace Projection (NSP) metric they developed uses no information about the target or the detector. They do, however, form a binary hypothesis test based on the eigenvalues of the whitened image covariance matrix.

The algorithm begins by estimating the image covariance matrix from the data. The inverse of the covariance matrix is decomposed such that

$$\mathbf{C}^{-1} = \mathbf{D}\mathbf{E}\mathbf{D} \quad (22)$$

where \mathbf{D} is a diagonal matrix created from the square root of the diagonal elements of \mathbf{C}^{-1} and \mathbf{E} is a matrix of correlation coefficients of \mathbf{C}^{-1} . Using this decomposition, the whitening matrix is defined as

$$\mathbf{W} = \mathbf{D}^{-1}. \quad (23)$$

Using (23), the image covariance matrix is whitened such that

$$\mathbf{C}_w = \mathbf{W}\mathbf{C}\mathbf{W}. \quad (24)$$

The whitening is performed to reduce the correlations inherent between spectral bands. The whitened matrix is analyzed using Principal Component Analysis (PCA) to extract the eigenvalues for the binary hypothesis test. The hypotheses are

$$\begin{aligned} H_0 : \lambda_i &= 1 \\ H_1 : \lambda_i &> 1 \end{aligned} \quad (25)$$

for each i^{th} eigenvalue. The likelihood function for the null hypothesis is simplified to

$$p_o(\lambda_i) \cong N(1, \frac{2}{N}). \quad (26)$$

Using (26), a threshold can be calculated for a given false alarm probability. Because (26) is independent of the index i , the same threshold can be applied to all eigenvalues. Using this information, the number of endmembers can be found using

$$\hat{m} = \max_m \left(\frac{(\lambda_m - 1)\sqrt{N}}{\sqrt{2}} \geq \Phi^{-1}(1 - p) \right) \quad (27)$$

where $\Phi^{-1}(1-p)$ is the inverse of the standard normal cumulative density function (cdf) evaluated at probability $1-p$. From [19], they recommend a value of 0.001 for p .

4.3.2.2. Thai/Healey Metric

This metric was developed as an aside in Thai and Healey's invariant subpixel detection paper [109]. The paper is another variant of the AMSD algorithm where the

target subspace is created using Healey and Slater's invariant method [45]. Thai applied this invariant method to subpixel detection and independently derived the AMSD algorithm [71]. Unlike Manolakis, Siracusa, and Shaw [71] who depended on the energy estimate described earlier, Thai and Healey designed a new metric to choose the dimension of their background subspace.

The basic idea is to find the number of endmembers that maximize target detection while minimizing the background. To accomplish this, they created a ratio of AMSD statistics such that

$$\hat{m} = \max_{m \in \mathbf{M}} \left(\frac{\delta_{AMSD}(\bar{\mathbf{S}})}{\delta_{AMSD}(\boldsymbol{\mu})} \right) \quad (28)$$

where $\delta_{AMSD}(\mathbf{x})$ is the AMSD statistic given in Chapter 5, $\bar{\mathbf{S}}$ is the mean of the target signatures, and $\boldsymbol{\mu}$ is the adjusted spectral mean of the image. The adjusted spectral mean is calculated from all the pixels in the image except those whose matched filter score is near one. The set of \mathbf{M} restricts the values of m based on the mean squared error between the original image and the PCA decomposition of the image. Thus, at least m_1 eigenvectors are always used, but not more than m_2 eigenvectors. The reasoning is that the number of eigenvectors that make up the background must be numerous enough to minimize the mean squared error, but not so numerous that the eigenvectors are pure "noise." No discussion is provided on how to derive these limits, or the threshold used in the matched filter.

4.3.2.3. AMSD MDL Metric

The first proposed metric fuses the ideas from the AMSD detector and MDL criterion. The original MDL equation in (19) can be formed for any likelihood. The method used by Du and Chang in their paper assumed that the HSI data could be fully

modeled by the image covariance matrix following Wax and Kailath's work. Unfortunately, this approach is not applicable to hyperspectral analysis as previously discussed in [19]

Instead of the image covariance matrix, we propose using the AMSD likelihood directly in the MDL criterion. This would match the criterion to the specific detector and all of its implicit assumptions. However, with any detector, there are two likelihoods: one for the null hypothesis and one for the alternate hypothesis. For this criterion, we use the alternate hypothesis which includes the target signature(s). The reasoning is the alternate hypothesis includes information about the target signature as well as the detector. Therefore, combining the MDL criterion with the alternate AMSD likelihood is

$$\hat{m} = \min_m \left(L \sum_{i=1}^N \left(\mathbf{x}_i^T (\mathbf{I} - \mathbf{E}_m (\mathbf{E}_m^T \mathbf{E}_m)^{-1} \mathbf{E}_m^T) \mathbf{x}_i \right) + m(N+1) \log N \right) \quad (29)$$

where \mathbf{E}_m is the concatenation of the target and m background signatures, L is the number of spectral bands, and N is the number of pixels in the image.

4.3.2.4. Subpixel Dimensionality Metric

In the MDL AMSD criterion, the idea was to identify the number of endmembers that minimized the likelihood of the denominator. This is only part of the optimization problem however. Ideally, the number of endmembers should also maximize the numerator. Interestingly, this is the same optimization done in detection theory; so, we can use the detector directly to estimate the number of endmembers. Following the approach of Thai and Healey [109], the subpixel dimensionality metric is

$$\hat{m} = \max_{m \in M} (\delta_{AMSD}(a\bar{\mathbf{S}} + (1-a)\boldsymbol{\mu})) \quad (30)$$

where the dimensionality is chosen by maximizing the AMSD detection score for a simulated pixel that is a linear combination of the desired target spectra and spectral mean from the image. The abundance a is calculated a-priori given the size of the target and the size of the pixels based on the altitude and the sensor's field of view parameters.

This approach has a number of advantages. First, like Thai and Healey's method, the metric can be quickly calculated for all numbers of endmembers. Second, the statistic directly uses the detector accounting for application dependencies unlike the intrinsic dimensionality metrics. Third, the metric chooses the number of endmembers based on both the predicted size and spectral characteristics of the target.

4.4. Experimental Results

Endmember extraction algorithms have been compared in a number of papers [81],[82],[116], but little experimentation has been performed on the impact of background dimensionality on subpixel target detection performance. This section compares the different methods of background dimensionality estimation and their impact on subpixel target detection. The goal is to identify which methods provide good dimensionality estimates for subpixel detection and under what conditions.

The experiments are broken into two parts: individual image results and ROC results. The individual image results present P_d and P_{fa} results for each image and target type. The ROC performance provides results across all images including those that do not contain targets. All the experiments use the Sensor X data for Images 1

through 6 and Targets 1 through 4. As mentioned in Chapter 2, Targets 1 and 2 are relatively easy to identify. Target 3 is more difficult because of the inherent variability in the spectral signature. Target 4 is very difficult to detect due to its low reflectance.

The detector used for these experiments is the AMSD detector described in Chapter 5. This is a standard structured subpixel detector in the literature that uses the eigenvectors of the image correlation matrix as the background endmembers. This type of detector allows us to apply all of the background dimensionality estimates on similar background information (image covariance or image correlation matrix).

4.4.1. Individual Image Results

Tables 5 through 8 provide the results of the individual image experiments for Targets 1 through 4 respectively. In each table, the number of endmembers (m), the P_d , and the number of false alarms (FA) are provided for each of the background dimension estimates described in Section 4.3. The ideal case is also provided in the last column. This case was found using the known ground truth to find the number of endmembers providing the highest P_d while minimizing the number of false alarms. Each table includes only the images in which targets are present.

The results show some intriguing results. First, the energy metric does not perform well as expected. For this implementation, we required 99.9% of the energy be obtained leading to background estimates of 2 to 3 endmembers. Unfortunately in radiance space, these first few eigenvectors comprise most of the environmental effects. This has the effect of providing little separation between target and background for all target types. The P_d is typically low with high false alarm rates.

The impact of this finding shows that papers using this metric [60],[71] are biasing their results against the AMSD detector.

Table 5: Comparison of Dimensionality Estimates for Target 1

Metric	Image	Energy	EIF	NSP	Thai/ Healey	AMSD MDL	SDD	Ideal
m	1	3	103	97	12	11	121	68
	4	2	105	97	39	11	131	125
P_d	1	0.15	1.00	1.00	0.95	1.00	1.00	1.00
	4	0.05	1.00	1.00	1.00	1.00	1.00	1.00
FA	1	36	0	0	280	334	2	0
	4	66	0	0	10	383	0	0

Table 6: Comparison of Dimensionality Estimates for Target 2

Metric	Image	Energy	EIF	NSP	Thai/ Healey	AMSD MDL	SDD	Ideal
m	1	3	103	97	5	11	111	11
	4	2	105	97	6	9	133	7
P_d	1	0.74	1.00	1.00	0.71	1.00	1.00	1.00
	4	0.90	1.00	1.00	1.00	1.00	1.00	1.00
FA	1	175	253	130	212	27	360	27
	4	271	34	17	5	43	92	0

The Thai/Healey and AMSD MDL metrics perform poorly as well. This is a surprising result as these metrics use knowledge of the target signature and detector type to estimate the background dimension. The Thai/Healey metric degrades significantly as the targets become more difficult to detect. Even on the simpler targets, the P_d is less than the other methods with higher false alarm densities. The reason this occurs is because the targets are truly subpixel, but the metric assumes a full pixel target. This causes a mismatch between what is being estimated and what is present in the data. We would expect the metric to perform well on full-pixel targets even though it was developed for subpixel target applications.

Table 7: Comparison of Dimensionality Estimates for Target 3

Metric	Image	Energy	EIF	NSP	Thai/ Healey	AMSD MDL	SDD	Ideal
m	2	3	102	89	18	12	61	31
	3	3	104	89	1	4	82	108
	5	3	106	99	6	11	104	79
	6	3	106	98	1	11	103	117
P_d	2	0.00	1.00	1.00	0.92	0.75	0.92	1.00
	3	0.56	1.00	1.00	0.12	0.68	1.00	1.00
	5	0.20	0.93	0.93	0.67	0.73	0.93	1.00
	6	0.57	0.96	0.93	0.07	0.96	0.96	1.00
FA	2	332	108	711	58	340	248	21
	3	492	3	9	44	112	17	0
	5	13	28	107	35	107	237	339
	6	139	6	5	55	81	4	0

The AMSD MDL criterion also degrades significantly as the targets become more difficult to detect. For Target 1, the criterion does find all the targets, but also provides the highest false alarm numbers. For Target 2, AMSD MDL outperforms the other metrics obtaining estimates close to the ideal. On the last two targets, the AMSD MDL estimate degrades losing significant P_d and obtaining large false alarm densities. The estimates vary because the metric is only treating the denominator of the AMSD statistic without reference to the effect of the numerator. In Target 2, this is not a significant problem, but for all other target types, the numerator decreases as quickly as the denominator causing the metric to erroneously pick the wrong number of endmembers.

The last three metrics (EIF, NSP, and SDD) perform well. The EIF criterion does well without any information about detector type or target signature. This is an interesting result as the other two methods are virtual dimensionality statistics. However, the EIF criterion was developed for identifying the number of spectral

signatures in chemical spectroscopy. This idea seems to have merit when applied to optical spectroscopy even with lack of target and detector knowledge.

Table 8: Comparison of Dimensionality Estimates for Target 4

Metric	Image	Energy	EIF	NSP	Thai/ Healey	AMSD MDL	SDD	Ideal
m	2	3	102	89	4	8	88	6
	3	3	104	89	3	13	83	6
	5	3	106	99	6	10	87	134
	6	3	106	98	1	11	85	70
P_d	2	0.33	0.00	0.00	0.11	1.00	0.00	1.00
	3	0.00	0.22	0.43	0.00	0.00	0.09	0.52
	5	0.42	0.25	0.17	0.00	0.17	0.00	0.50
	6	0.04	0.04	0.12	0.00	0.04	0.00	0.28
FA	2	279	786	766	186	234	721	230
	3	444	789	581	444	230	677	472
	5	109	649	623	429	450	661	898
	6	310	516	493	20	210	755	673

The NSP algorithm which was developed for HSI data performs well. The estimate provides some of the lowest false alarm densities for Targets 1 and 2 while maintaining 100% P_d . As with the other methods, NSP breaks down as the targets become more difficult; however, it does not degrade as fast as energy or AMSD MDL. NSP, in fact, maintains the highest number of target detections on Target 4.

The final algorithm is the proposed SDD metric. This metric performs similarly to the EIF and NSP metrics. As expected, the performance of this metric is directly linked to the detector performance. When the targets become more difficult for the detector to find, this metric degrades as well. The unfortunate outcome of this result is that it provides some of the worst performance on Target 4, but some of the best performance on Target 1.

The overall results of these experiments are mixed. The energy metric is not desirable due to its poor performance across images. The AMSD MDL metric is not desirable due to its variable performance that is uncorrelated with the difficulty of the target type. Thai and Healey's metric which does not account for the subpixel nature of the target provides poor estimates as well. The EIF, NSP, and SDD metrics perform well and degrade gracefully as the target becomes more difficult to detect.

4.4.2. ROC Results

The results from the first experiment show the EIF, NSP, and SDD estimates perform similarly well when applied to images with targets. To see if any separation exists between these methods, it is interesting to look at cases where images that do not contain targets are used. To measure the effect of such images, we use ROC curves.

ROC curves show the average performance of the detector across all images. Ideally, the number of endmembers used should help suppress the background pixels into the same range of detection scores. This allows the ROC curve to apply the same threshold across each image and get similar results. When the background is not confined to the same range of detection scores, the background detection scores on one image may actually be higher than the target detection scores on another image. In such cases, the inconsistency of the detection scores will negatively impact ROC performance. Thus, the impact of images without targets can be assessed on the overall detector performance.

Figures 21 through 24 provide the ROC curves for Targets 1 through 4 respectively. In each figure, there are seven curves. The first curve represents the ideal based on ground truth information obtained with the imagery. The other six

curves represent the estimation algorithms defined in Section 4.3. Similar to the previous experiment, the ideal number of endmembers was selected as those that maximized P_d while minimizing the number of false alarms. In the images without targets, the ideal was chosen as the number of endmembers that suppressed the detection scores into ranges that were similar to the other images.

As expected given the earlier experimental results, the energy, Thai/Healey, and AMSD MDL criteria did not perform well. While these results do not provide good performance, they do highlight the need for good background dimension estimates. The interesting exception to this rule is the AMSD MDL curve for Target 4. For this target, the AMSD MDL curve is one of the best, but this is most likely a coincidence as the estimate simply favors lower numbers of endmembers.

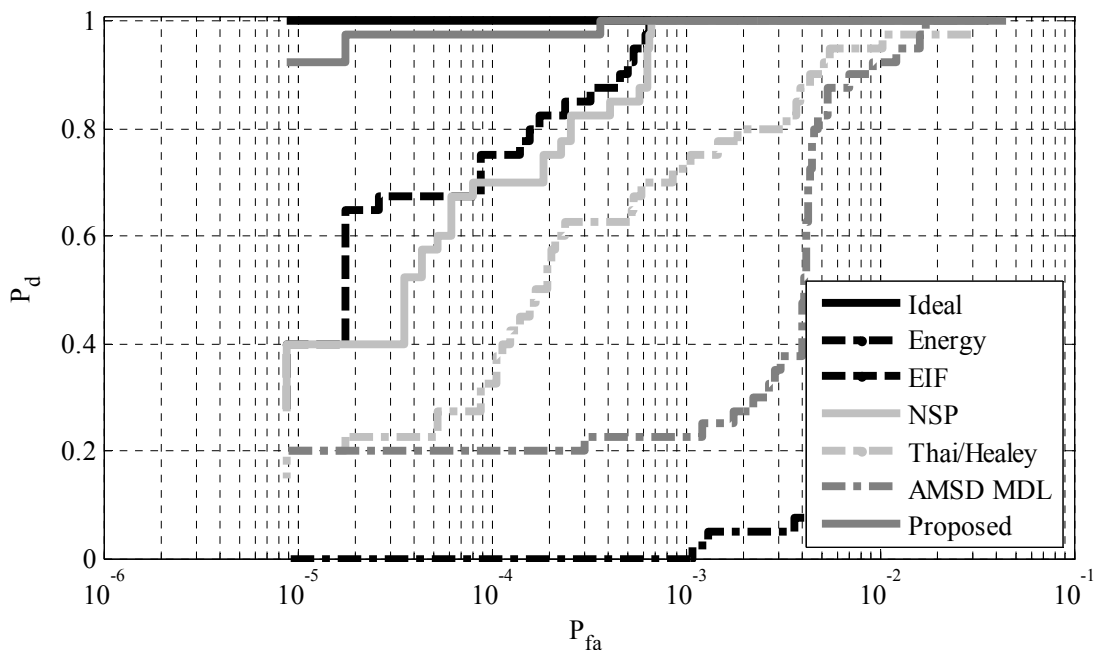


Figure 21: Comparison of Background Dimension Estimates for Target 1

The interesting results occur with the EIF, NSP, and SDD methods. In the first set of experiments, these algorithms perform nearly equally well on the different targets and images. In these ROC experiments however, the algorithms respond

differently. The EIF and NSP methods performance is best with Targets 2 and 3. These targets are easy to moderately difficult to detect. For Target 1, the methods perform significantly worse than the ideal case. For Target 4, the NSP method performs nearly the best although this is again significantly less than the ideal. Nevertheless, both algorithms are consistently some of the best methods for background dimension estimation.

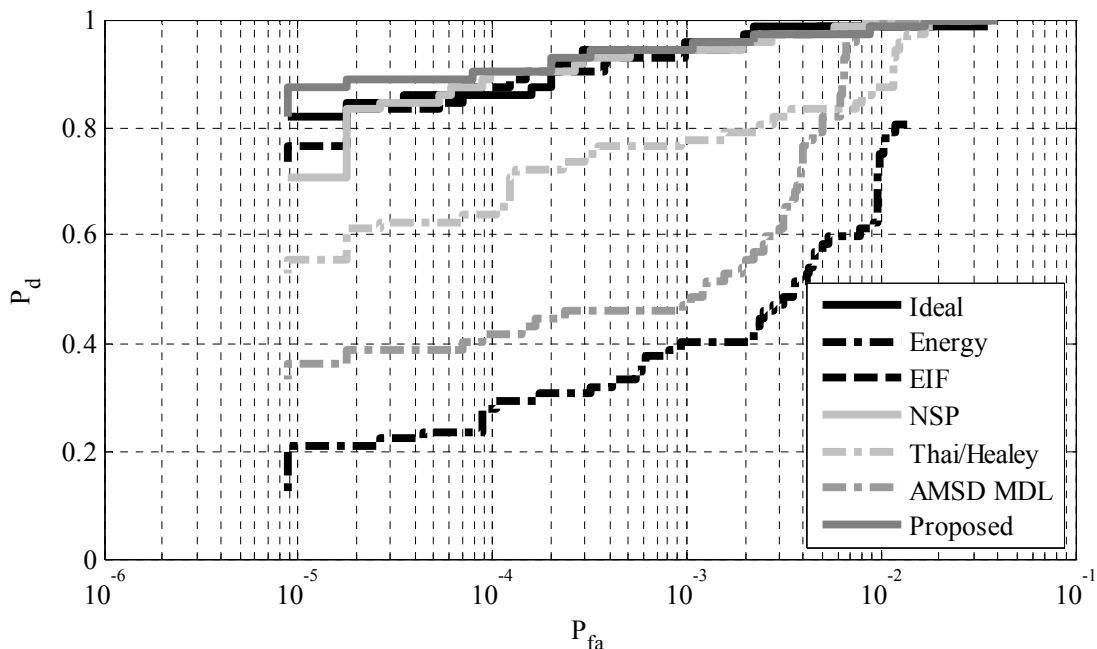


Figure 22: Comparison of Background Dimension Estimates for Target 2

The SDD method demonstrates excellent performance when the target is easy and degrades as the targets become more difficult. This performance is expected given the method is based directly on the performance of the detector using a simulated subpixel target. For Target 1, the SDD method is nearly ideal and substantially better than any other method tested. For Target 2, the method matches the ideal case although the EIF and NSP methods have similar performance. For Target 3, the SDD method degrades slightly as this target is more difficult to detect

due to the spectral variability of the target. On this target, the EIF and NSP methods have a slight advantage. For Target 4 however, the SDD method performs poorly because the detector has difficulty finding such a weak target.

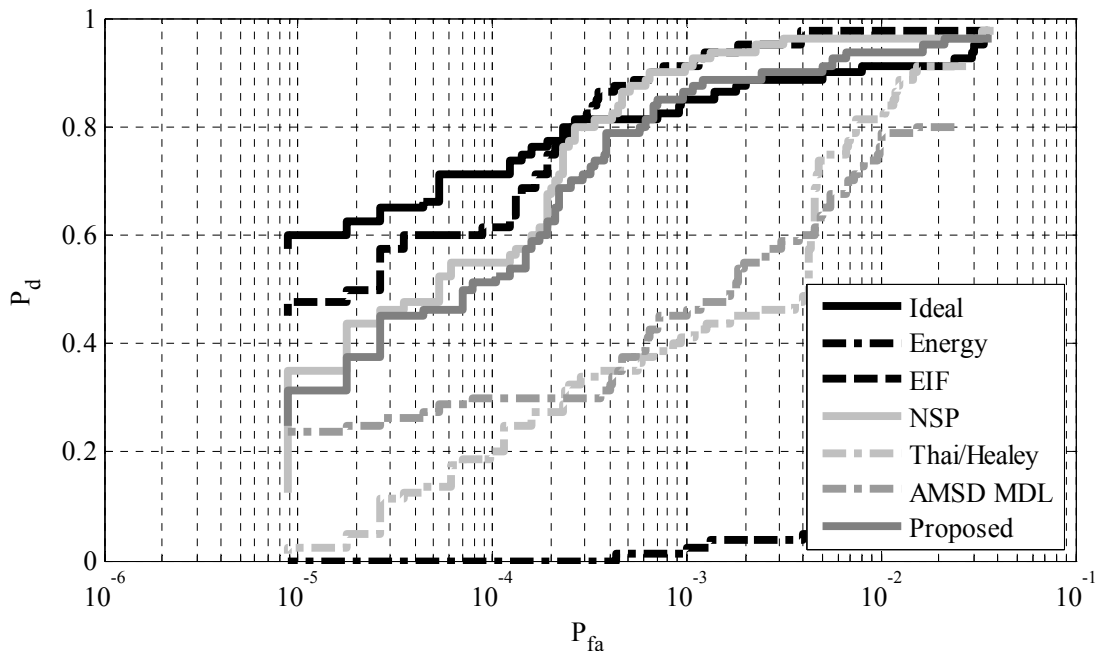


Figure 23: Comparison of Background Dimension Estimates for Target 3

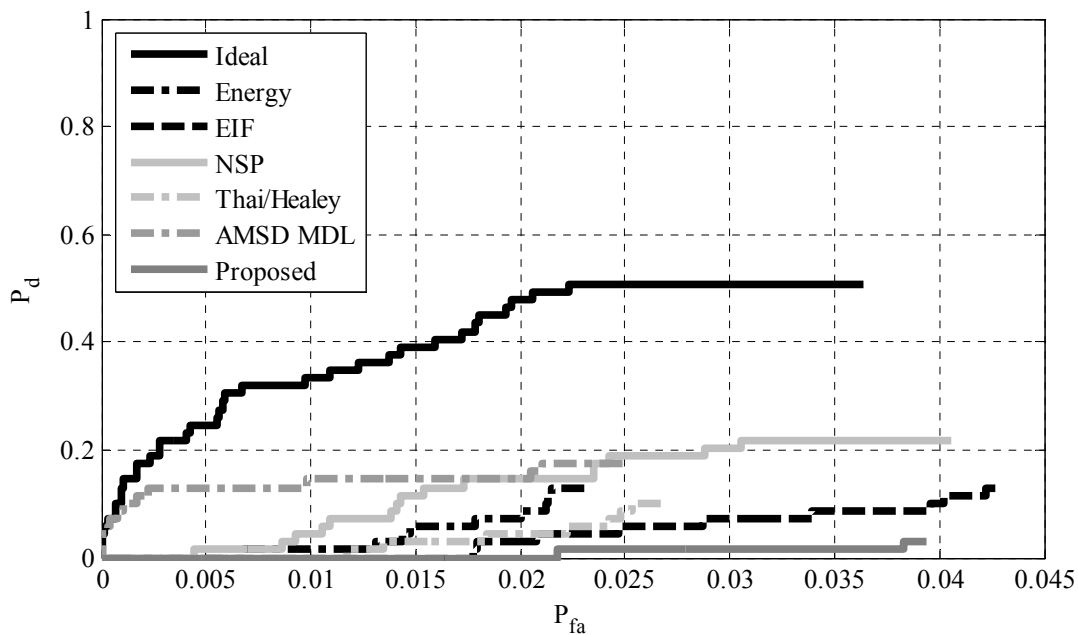


Figure 24: Comparison of Background Dimension Estimates for Target 4

4.4.3. *Conclusions*

The results of these experiments show that the SDD method has an advantage over the other estimates for all detectable targets. Since the method is based directly on the subpixel detector performance, this result is expected. The EIF and NSP methods are close competitors. These methods show good separation in both the single image and ROC experiments. Since these algorithms were intentionally designed for HSI data, these results are consistent with theory.

The energy, Thai/Healey, and AMSD MDL methods are not good indicators of the background dimension. Energy is the worst indicator although it has been used in numerous papers. The Thai/Healey method does not perform well despite being designed for subpixel processing using the AMSD algorithm. This can most likely be traced to the fact that Thai and Healey used mostly targets that were not subpixel in their paper. For full pixel targets, the method should work well. Unfortunately, the AMSD MDL method did not perform well because it only uses the denominator of the AMSD detector to make its estimate.

4.5. *Summary*

The estimation of the number of background endmembers for subpixel detection remains a challenging problem. Our work has shown that improvements can be made over the current methods, but these improvements are directly linked to the performance of the detector and the strength of the target signature. In cases where the target signature is well characterized and significantly different from the background, the SDD method we proposed works very well followed closely by the EIF and NSP methods. As the target becomes weaker (or the background becomes more complex), all of the methods degrade.

Further research should be continued to identify better ways to estimate the background dimension. The results clearly show the loss of performance when the background is not correctly identified. Such performance can be significant – especially in the case of weak targets like Target 4. The other direction is to develop detection algorithms that are partially invariant to the number of background endmembers. Such algorithms would show minimal loss in subpixel detection performance due to minor errors in background dimension estimation.

Chapter 5: Physics-Based Hybrid Detectors

A number of different methods have been proposed to address subpixel detection. One of the earliest methods uses array processing techniques to nullify the background signatures as one would nullify an interfering signature when performing beamforming. The Orthogonal Subspace Projection (OSP) [41] and Constrained Energy Minimization (CEM) [20] algorithms are examples of such methods. In order to implement these detectors, the authors assume the noise to be a zero-mean multivariate normal distribution with covariance matrix $\sigma^2 \mathbf{I}$. The idea behind this algorithm is that the background can be fully characterized by endmembers and that the remaining noise will meet the aforementioned $\sigma^2 \mathbf{I}$ assumption.

Another approach uses the linear mixing model to directly estimate the abundance values and use the estimated target abundances for detection purposes. Two examples of this approach are the Non-Negativity Constrained Least Squares [20] and Fully Constrained Least Squares algorithms [46]. These methods can be considered physics-based methods since they attempt to address all of the phenomenological constraints in the linear mixing model. Others have also incorporated the constraint of a full covariance matrix into these methods which can be considered as the first use of a semi-structured approach [86]. Incorporation of covariance information addresses the fact that most of the spectral bands in HSI data are highly correlated. The estimated covariance is used for designing a whitening transform that decorrelates the bands making the HSI data fit the aforementioned assumption of $\sigma^2 \mathbf{I}$. These physics-based methods perform well for both unsupervised estimation of background endmembers and the calculation of the corresponding

abundances; however, they do not provide a statistical hypothesis test – they only provide an estimate of the target abundance.

To develop such a statistical test, a set of hypotheses must be generated to differentiate those pixels containing targets of interest from those pixels that exclusively contain background spectra. The set of hypotheses are

$$\begin{aligned} H_0 : \mathbf{x} &= \mathbf{B}\mathbf{a}_{b,0} + \mathbf{n} \\ H_1 : \mathbf{x} &= \mathbf{S}\mathbf{a}_s + \mathbf{B}\mathbf{a}_{b,1} + \mathbf{n} \end{aligned} \quad (31)$$

where \mathbf{x} is the pixel under test, \mathbf{B} is a $L \times Q$ matrix representing background endmembers, $\mathbf{a}_{b,0}$ and $\mathbf{a}_{b,1}$ are the abundances of the background endmembers under each hypothesis, \mathbf{S} is a $L \times P$ matrix representing target endmembers, \mathbf{a}_s are the abundances of the targets, and \mathbf{n} is a noise model typically assumed to be a zero-mean multivariate normal distribution.

Using this set of hypotheses, a set of detectors has been developed based on structured and unstructured backgrounds. A good example of a structured background detector is the Adaptive Matched Subspace Detector (AMSD) [71]. The AMSD algorithm models the background using the linear mixing model with endmembers and abundances. This statement is a misnomer however since the endmembers in the AMSD algorithm have no physical meaning. Instead the endmembers are the eigenvectors of the image correlation matrix. Thus the abundances are no longer measurements of area. They are simply magnitudes along the eigenvector directions which in general do not satisfy the non-negativity and sum-to-one constraints. So although the linear mixing model is used as the basis for AMSD, all physical considerations are ignored in favor of a purely statistical approach. While AMSD has

shown good performance, research has shown that a purely structured background model does not fully represent the background in real-world HSI data [71].

An example of an unstructured detector is the Adaptive Cosine/Coherent Estimate (ACE) [58]. The ACE algorithm assumes no background signatures opting instead for modeling the background as a multivariate normal distribution. While this removes the need to extract and identify the proper number of background endmembers, it also removes the physical constraints of the linear mixing model. Despite this seemingly simple background model, the ACE detector is one of the more powerful subpixel detectors available for HSI data [70]. Unfortunately, research has shown that an unstructured detector which uses the multivariate normal distribution is not a good model of backgrounds in hyperspectral imagery [103].

Another algorithm that uses the hypotheses in (31) is the Constrained Signal Detector (CSD) [49]. This algorithm was the one of the first to use some of the physical constraints of the linear mixing model within a statistical hypothesis test. The algorithm included the sum-to-one constraint on the abundances, but only required the target abundance to be non-negative arguing that proper estimation of the background abundances was not required for detection purposes. The algorithm was also designed assuming that the noise was zero-mean multivariate normal distribution with covariance $\sigma^2 \mathbf{I}$. These assumptions made the algorithm very fast, but still do not account for all of the physical constraints in the linear mixing model or a full covariance for the background noise distribution.

Therefore, we present two new hybrid subpixel detectors based on modeling the background using a physically meaningful linear mixing model within a statistical

hypothesis test. The idea is that the physically-based endmembers and abundances will account for the known physics of the problem while the statistical distribution accounts for unknown quantities due to such phenomena as nonlinear mixing effects and sensor noise. Our hypothesis is that the hybrid detectors which model the background both physically and statistically will provide improved performance over their purely statistical counterparts AMSD and ACE. Section 5.1 describes the FCLS, AMSD, and ACE algorithms that form the basis for our hybrid detectors. Section 5.2 describes the two proposed hybrid detectors. Section 5.3 details the experiments used to test our hypothesis. Section 5.4 presents the results of the experiments showing the hybrid detectors excel in three areas: endmember insensitivity, target/background separation on an image by image basis, and improved ROC performance over multiple images. Section 5.5 summarizes the results and identifies future research directions.

5.1. Current Subpixel Algorithms

This section details the FCLS, AMSD, and ACE algorithms. These algorithms are the foundation on which we derive the hybrid detectors. The FCLS algorithm provides a method to incorporate the sum-to-one and non-negativity constraints on the abundances. The AMSD algorithm provides a detector based on a structured background that uses endmembers to define the background \mathbf{B} . The ACE algorithm provides a detector based on an unstructured background (i.e., a background modeled by a statistical distribution instead of endmembers).

5.1.1. Fully Constrained Least Squares (FCLS)

The FCLS algorithm directly estimates the abundances in (31). While other algorithms have been developed that handle both the non-negativity and sum-to-one

constraints [4][8][98], these algorithms tend to be computationally intense as the number of endmembers increase. The FCLS algorithm meets both abundance constraints as well, but in an efficient manner that is optimal in terms of least squares error (LSE) [46]. Because of these reasons, we chose to use it in our algorithms. Unfortunately, FCLS does not allow a closed-form mathematical solution due to the non-negativity constraints. Instead, a numerical solution is required.

To calculate the FCLS solution, we begin with the non-negativity constraints. The idea is to minimize the LSE by estimating the non-negative abundance values. Mathematically this is expressed as

$$\min_{\mathbf{a}} (\mathbf{x} - \mathbf{E}\mathbf{a})^T (\mathbf{x} - \mathbf{E}\mathbf{a}), a_i \geq 0 \forall i \quad (32)$$

where \mathbf{E} is the concatenation of the target \mathbf{S} and background \mathbf{B} signatures. Using Lagrange multipliers, a Lagrangian J is defined such that

$$J = \frac{1}{2} (\mathbf{x} - \mathbf{E}\mathbf{a})^T (\mathbf{x} - \mathbf{E}\mathbf{a}) + \boldsymbol{\lambda}^T (\mathbf{a} - \mathbf{c}), \quad (33)$$

where $\mathbf{a} = \mathbf{c}$, and each member of the unknown constant $M \times I$ vector \mathbf{c} is non-negative to enforce the non-negativity constraint. This construction allows the use of Lagrange multipliers because the non-negativity constraints have been substituted by equality constraints with the unknown vector \mathbf{c} . To calculate the estimate of \mathbf{a} , we take the partial derivative of J with respect to \mathbf{a} to obtain

$$\left. \frac{\partial J}{\partial \mathbf{a}} \right|_{\hat{\mathbf{a}}} = \mathbf{E}^T \mathbf{E} \hat{\mathbf{a}} - \mathbf{E}^T \mathbf{x} + \boldsymbol{\lambda} = 0. \quad (34)$$

Equation (34) contains two unknowns: the abundance estimates and the Lagrange multipliers. Solving for these unknown results in

$$\hat{\mathbf{a}} = (\mathbf{E}^T \mathbf{E})^{-1} \mathbf{E}^T \mathbf{x} - (\mathbf{E}^T \mathbf{E})^{-1} \boldsymbol{\lambda} \quad (35)$$

and

$$\boldsymbol{\lambda} = \mathbf{E}^T (\mathbf{x} - \mathbf{E}\hat{\mathbf{a}}). \quad (36)$$

Iterating through (35) and (36) provides the numerical solution for the non-negativity constraints. To begin this iterative method, we set all the Lagrange multipliers to zero and calculate the abundance using (35). Note that this initial calculation is the unconstrained least squares solution for the abundance values. From this solution, we identify those abundance values that are greater than zero and place them in the passive set P . The remaining non-positive abundance values are placed in the active set R . Equations (35) and (36) are iterated until all Lagrange multipliers in the passive set are zero and all Lagrange multipliers in the active set are either zero or negative. At this point, the Kuhn-Tucker conditions have been met and an optimal solution for the abundance values has been found.

One may note that this solution only accounts for the non-negativity constraints of (1). To handle the sum-to-one constraints, an easy modification of the aforementioned algorithm was developed to retain the optimality guaranteed under the Kuhn-Tucker conditions for numerical optimization on a finite computing machine [42]. In the modification, the endmember matrix and pixel signatures are extended such that

$$\tilde{\mathbf{E}} = \begin{bmatrix} \delta \mathbf{E} \\ \mathbf{1}^T \end{bmatrix} \quad (37)$$

is the new endmember matrix and

$$\tilde{\mathbf{x}} = \begin{bmatrix} \delta \mathbf{x} \\ 1 \end{bmatrix} \quad (38)$$

is the new pixel signature where δ is a small number (typically 1×10^{-5}). The δ variable controls how tightly the solution will sum to one so that smaller values provide a better solution, but may need longer convergence time. The new endmember matrix and pixel signature are then used in (35) and (36) to obtain an abundance solution that meets both the non-negativity and sum-to-one constraints.

5.1.2. Adaptive Matched Subspace Detector (AMSD)

While the FCLS algorithm provides an elegant solution to calculating the abundance values in the linear mixing model, the algorithm does not provide a statistical hypothesis test to differentiate between a pixel that contains a target and a pixel that contains only the background. The AMSD algorithm provides such a statistical test using a Generalized Likelihood Ratio Test (GLRT) [71]; however, the non-negativity and sum-to-one constraints on the abundance estimates are in general not satisfied. Thus, the AMSD approach leads to a closed-form solution with CFAR optimality, but has to sacrifice the physical constraints on the abundance estimates.

Since the AMSD algorithm is based on a GLRT, we can use the model in (31) assuming that the noise model is a zero-mean normal distribution with covariance matrix $\sigma^2 \mathbf{I}$. Therefore, the AMSD hypotheses are

$$\begin{aligned} H_0 : \mathbf{x} &\sim N(\mathbf{B}\mathbf{a}_{b,0}, \sigma_0^2 \mathbf{I}) \\ H_1 : \mathbf{x} &\sim N(\mathbf{S}\mathbf{a}_s + \mathbf{B}\mathbf{a}_{b,1}, \sigma_1^2 \mathbf{I}) \end{aligned} \quad (39)$$

Under these assumptions, we can calculate the remaining unknown parameters using Maximum Likelihood Estimation (MLE) techniques. To do this, we calculate the likelihood equation for the null hypothesis as

$$L(\mathbf{x} | H_0) = (2\pi\sigma_0^2)^{-\frac{L}{2}} \exp\left\{-\frac{1}{2\sigma_0^2} (\mathbf{x} - \mathbf{B}\mathbf{a}_{b,0})^T (\mathbf{x} - \mathbf{B}\mathbf{a}_{b,0})\right\}. \quad (40)$$

Taking the derivative of the logarithm of (40) with respect to each of the unknown parameters and setting them equal to zero allows us to arrive at the MLE abundance estimate

$$\hat{\mathbf{a}}_{b,0} = (\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T \mathbf{x} \quad (41)$$

and the MLE noise variance estimate

$$\hat{\sigma}_0^2 = \frac{1}{L} (\mathbf{x} - \mathbf{B} \hat{\mathbf{a}}_{b,0})^T (\mathbf{x} - \mathbf{B} \hat{\mathbf{a}}_{b,0}). \quad (42)$$

Substituting (41) and (42) back into (40) provides the generalized likelihood equation under the null hypothesis

$$f_0 = \left(\frac{2\pi}{L} \mathbf{x}^T (\mathbf{I} - \mathbf{B}(\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T) \mathbf{x} \right)^{-\frac{L}{2}} \exp\left(-\frac{L}{2}\right). \quad (43)$$

Similarly, the same can be done for the alternative hypothesis to arrive at

$$f_1 = \left(\frac{2\pi}{L} \mathbf{x}^T (\mathbf{I} - \mathbf{E}(\mathbf{E}^T \mathbf{E})^{-1} \mathbf{E}^T) \mathbf{x} \right)^{-\frac{L}{2}} \exp\left(-\frac{L}{2}\right) \quad (44)$$

where \mathbf{E} is again defined as the concatenation of the target and background signatures.

Having calculated the likelihoods for each hypothesis and using some simple algebra, the GLRT takes the ratio of the two likelihoods to calculate the following detection statistic

$$\frac{f_1}{f_0} \Rightarrow \frac{\mathbf{x}^T (\mathbf{I} - \mathbf{B}(\mathbf{B}^T \mathbf{B})^{-1} \mathbf{B}^T) \mathbf{x}}{\mathbf{x}^T (\mathbf{I} - \mathbf{E}(\mathbf{E}^T \mathbf{E})^{-1} \mathbf{E}^T) \mathbf{x}} = \frac{\mathbf{x}^T \mathbf{P}_B^\perp \mathbf{x}}{\mathbf{x}^T \mathbf{P}_Z^\perp \mathbf{x}}. \quad (45)$$

Since \mathbf{E} and \mathbf{B} are related, it is difficult to identify the distribution of this detection statistic; so, a new detection statistic is created by subtracting one from (45) to obtain

$$D_{AMSD}(\mathbf{x}) = \frac{\mathbf{x}^T (\mathbf{P}_B^\perp - \mathbf{P}_Z^\perp) \mathbf{x}}{\mathbf{x}^T \mathbf{P}_Z^\perp \mathbf{x}}. \quad (46)$$

Applying this mapping does not change the outcome of the decision statistic, but it does allow the new statistic to be distributed as

$$D_{AMSD}(\mathbf{x}) \sim F_{P,L-P-Q} \left(\frac{\|\mathbf{P}_B^\perp \mathbf{S} \mathbf{a}_s\|^2}{\sigma^2} \right) [71]. \quad (47)$$

Under the null hypothesis ($\mathbf{S} = \mathbf{0}$ and hence the signal to interference plus noise ratio (SINR) term in the parentheses of (47) is equal to zero), the AMSD statistic is based on the parameters P , L , and Q independent of any estimates. Because of this, the AMSD statistic enjoys the CFAR property and should allow a single threshold to determine the false alarm rate. Of course, the single threshold only holds if the underlying data has a multivariate normal distribution.

5.1.3. Adaptive Cosine/Coherent Detector

The methods described earlier are detectors based on structured backgrounds. The ACE method uses a statistical distribution (namely the multivariate normal distribution) to model the background. Referring to (31), the ACE algorithm sets $\mathbf{B} = \mathbf{0}$ thus removing any structured background information. In this algorithm, the background is entirely modeled as a zero-mean Gaussian distribution with scaled covariance $\sigma^2 \mathbf{\Gamma}$ giving us the hypotheses

$$\begin{aligned} H_0 : \mathbf{x} &\sim N(0, \sigma_0^2 \mathbf{\Gamma}) \\ H_1 : \mathbf{x} &\sim N(\mathbf{S} \mathbf{a}_s, \sigma_1^2 \mathbf{\Gamma}). \end{aligned} \quad (48)$$

The scaling term σ^2 is interesting as this term is not typically found empirically. The term is necessary theoretically however to make the ACE detector scale-invariant as will be shown later in this section. Since \mathbf{B} does not exist in this

algorithm, the sum-to-one and non-negativity constraints of (1) cannot be met either as they require a background subspace. Despite these seemingly poor assumptions for hyperspectral data, the ACE detector is one of the more powerful subpixel detectors available [70].

For this derivation, we follow the work by Kelly [53] and Kraut and Scharf [56][57][58]. Besides the information we have in (48), we also assume that we have an independent data set Y such that

$$Y = \{\mathbf{y}_i | \mathbf{y}_i \sim N(0, \mathbf{\Gamma}), i = 1, \dots, N\}. \quad (49)$$

Combining (48) and (49) provides the joint likelihood equation under the null hypothesis

$$L(\mathbf{x}, Y | H_0) = (2\pi)^{-\frac{1}{2}L(N+1)} |\mathbf{\Gamma}|^{-\frac{1}{2}(N+1)} (\sigma_0^2)^{-\frac{1}{2}L} \times \exp\left\{-\frac{1}{2\sigma_0^2} \mathbf{x}^T \mathbf{\Gamma}^{-1} \mathbf{x} - \frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^T \mathbf{\Gamma}^{-1} \mathbf{y}_i\right\} \quad (50)$$

and the joint likelihood equation under the alternate hypothesis

$$L(\mathbf{x}, Y | H_1) = (2\pi)^{-\frac{1}{2}L(N+1)} |\mathbf{\Gamma}|^{-\frac{1}{2}(N+1)} (\sigma_1^2)^{-\frac{1}{2}L} \times \exp\left\{-\frac{1}{2\sigma_1^2} (\mathbf{x} - \mathbf{S}\mathbf{a}_s)^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{S}\mathbf{a}_s)\right\} \times \exp\left\{-\frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^T \mathbf{\Gamma}^{-1} \mathbf{y}_i\right\}. \quad (51)$$

If we assume that N is very large, the covariance estimate from these likelihoods can be simplified to

$$\hat{\mathbf{\Gamma}} = \sum_{i=1}^N \mathbf{y}_i \mathbf{y}_i^T \quad (52)$$

which is a standard assumption made in the literature. Note that under this assumption, the covariance under the null hypothesis and alternate hypothesis are equal and greatly simplifying the following mathematics.

Following the derivation of the covariance under each hypothesis using MLE, we obtain the abundance estimate as

$$\hat{\mathbf{a}}_s = (\mathbf{S}^T \mathbf{\Gamma}^{-1} \mathbf{S})^{-1} \mathbf{S}^T \mathbf{\Gamma}^{-1} \mathbf{x} \quad (53)$$

and the variance estimates under each hypothesis as

$$\hat{\sigma}_0^2 = \frac{1}{L} \mathbf{x}^T \mathbf{\Gamma}^{-1} \mathbf{x} \quad (54)$$

and

$$\hat{\sigma}_1^2 = \frac{1}{L} (\mathbf{x} - \mathbf{S} \hat{\mathbf{a}}_s)^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{S} \hat{\mathbf{a}}_s). \quad (55)$$

The estimates are substituted back into the original likelihood equations in (50) and (51). The updated likelihoods are taken as a ratio to obtain the GLRT as was done in the AMSD derivation. After some algebra and simplification, the ACE detector is

$$D_{ACE}(\mathbf{x}) = \frac{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{S} (\mathbf{S}^T \mathbf{\Gamma}^{-1} \mathbf{S})^{-1} \mathbf{S}^T \mathbf{\Gamma}^{-1} \mathbf{x}}{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{x}}. \quad (56)$$

This is a CFAR detector and has the following distribution under the null hypothesis

$$D_{ACE}(\mathbf{x}) \sim \text{Beta}\left(\frac{P}{2}, \frac{L-P}{2}\right) \quad (57)$$

where L is the number of spectral bands and P is the number of target signatures [57]. Therefore, the ACE statistic is based only on the parameters P and L independent of any estimates. Because of this, the ACE statistic also enjoys the CFAR property and should allow a single threshold to determine the false alarm rate. Again, the single

threshold only holds if the underlying data is a multivariate normal distribution (or any distribution in the family of elliptically contoured distributions) [57].

5.2. *Hybrid Detectors*

Using the derivations and ideas in the previous section, we present two hybrid subpixel detectors that incorporate the HSI physical constraints directly into the detector derivation. The first detector uses a structured background and is similar to AMSD. The second detector uses an unstructured background and is similar to ACE.

5.2.1. *Hybrid Structured Detector*

The hybrid structured detector (HSD) approaches the solution to (31) using a structured background like AMSD, but using physically meaningful endmembers and replacing the abundance estimates with their FCLS counterparts. The HSD hypotheses are

$$\begin{aligned} H_0 : \mathbf{x} &\sim N(\mathbf{B}\mathbf{a}_{b,0}, \sigma_0^2 \mathbf{\Gamma}) \\ H_1 : \mathbf{x} &\sim N(\mathbf{S}\mathbf{a}_s + \mathbf{B}\mathbf{a}_{b,1}, \sigma_1^2 \mathbf{\Gamma}) \end{aligned} \quad (58)$$

Since this derivation includes a full covariance matrix, we follow a similar derivation to ACE incorporating the background subspace \mathbf{B} and its abundances \mathbf{a}_b as was done in AMSD. With this new information the likelihood equation under the null hypothesis is

$$\begin{aligned} L(\mathbf{x}, Y | H_0) &= (2\pi)^{-\frac{1}{2}L(N+1)} |\mathbf{\Gamma}|^{-\frac{1}{2}(N+1)} (\sigma_0^2)^{-\frac{1}{2}L} \times \\ &\exp \left\{ -\frac{(\mathbf{x} - \mathbf{B}\mathbf{a}_{b,0})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{B}\mathbf{a}_{b,0})}{2\sigma_0^2} - \frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^T \mathbf{\Gamma}^{-1} \mathbf{y}_i \right\} \end{aligned} \quad (59)$$

and the likelihood equation under the alternate hypothesis is

$$L(\mathbf{x}, Y | H_1) = (2\pi)^{-\frac{1}{2}L(N+1)} |\mathbf{\Gamma}|^{-\frac{1}{2}(N+1)} (\sigma_1^2)^{-\frac{1}{2}L} \times \exp \left\{ -\frac{(\mathbf{x} - \mathbf{Ea})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{Ea})}{2\sigma_1^2} - \frac{1}{2} \sum_{i=1}^N \mathbf{y}_i^T \mathbf{\Gamma}^{-1} \mathbf{y}_i \right\} \quad (60)$$

where $\mathbf{Ea} = \mathbf{Sa}_s + \mathbf{Ba}_{b,1}$.

The covariance estimate is the same as (52) given the assumption that N is large. Under this assumption, we obtain the variance estimates under each hypothesis as

$$\hat{\sigma}_0^2 = \frac{1}{L} (\mathbf{x} - \mathbf{Ba}_{b,0})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{Ba}_{b,0}) \quad (61)$$

and

$$\hat{\sigma}_1^2 = \frac{1}{L} (\mathbf{x} - \mathbf{Ea})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{Ea}). \quad (62)$$

Besides the covariance and variance estimates, the abundance estimates also need to be calculated. At this point instead of using the standard MLEs, we use a variant of the FCLS algorithm to estimate these parameters. Because of the covariance matrix, the variant of the FCLS algorithm attempts to minimize

$$\min_a (\mathbf{x} - \mathbf{Ea})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{Ea}), a_i \geq 0 \forall i \quad (63)$$

This update leads to a new Lagrangian J such that

$$J = \frac{1}{2} (\mathbf{x} - \mathbf{Ea})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{Ea}) + \lambda(\mathbf{a} - \mathbf{c}). \quad (64)$$

Therefore, the new equations that we iterate through to meet the Kuhn-Tucker conditions are

$$\hat{\mathbf{a}} = (\mathbf{E}^T \mathbf{\Gamma}^{-1} \mathbf{E})^{-1} \mathbf{E}^T \mathbf{\Gamma}^{-1} \mathbf{x} - (\mathbf{E}^T \mathbf{\Gamma}^{-1} \mathbf{E})^{-1} \lambda \quad (65)$$

and

$$\lambda = \mathbf{E}^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{E}\hat{\mathbf{a}}). \quad (66)$$

The rest of the algorithm proceeds as in Section 5.1.1 to obtain abundance estimates that incorporate the sum-to-one and non-negativity constraints with a full covariance matrix. While this prevents us from obtaining a closed-form solution for our detector, it enforces all of the known physical constraints.

All of the estimates are substituted back into the original likelihood equations in (59) and (60). The generalized likelihoods are taken as a ratio to obtain the GLRT as was done in the AMSD derivation. After some algebra and simplification, the HSD is

$$D_{HSD}(\mathbf{x}) = \frac{(\mathbf{x} - \mathbf{B}\hat{\mathbf{a}}_b)^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{B}\hat{\mathbf{a}}_b)}{(\mathbf{x} - \mathbf{E}\hat{\mathbf{a}})^T \mathbf{\Gamma}^{-1} (\mathbf{x} - \mathbf{E}\hat{\mathbf{a}})}. \quad (67)$$

The HSD algorithm is similar to our original hybrid detector [16] except for the inclusion of the full covariance matrix.

5.2.2. Hybrid Unstructured Detector

The Hybrid Unstructured Detector (HUD) models the background as a multivariate normal distribution similar to ACE. Since the ACE detector is already white, the HUD algorithm simply replaces the abundance estimates with their whitened FCLS counterparts. To accomplish this, we rewrite (56) such that

$$D_{ACE}(\mathbf{x}) = \frac{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{S} \mathbf{a}}{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{x}} \quad (68)$$

where the abundance estimate \mathbf{a} is taken from (53).

To form the HUD algorithm, we simply replace the abundance with its whitened FCLS counterpart. Therefore, the new detector is

$$D_{HUD}(\mathbf{x}) = \frac{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{S} \hat{\mathbf{a}}}{\mathbf{x}^T \hat{\mathbf{\Gamma}}^{-1} \mathbf{x}} \quad (69)$$

where the abundance estimate \mathbf{a} is taken from (65) after the Kuhn-Tucker conditions have been satisfied. Note that this solution still requires the extraction of endmembers to define the abundance estimate, but these endmembers are not directly used within the decision statistic. They only serve to provide a better estimate of the target abundance based on the physical constraints of the linear mixing model.

5.3. *Experimental Results*

Our hypothesis is that the hybrid detectors provide improved performance by taking advantage of the known physics of the linear mixing model within a statistical hypothesis test. To show whether this occurs or not in practice, we have implemented a number of experiments on hyperspectral imagery under real-world conditions. One of the major difficulties in doing such an analysis is being as unbiased as possible. This is a real concern when using real world hyperspectral data as many of the variables are simply out of our control. However, we can develop a series of tests that reduce this bias and provide meaningful results. We argue that these types of tests are more germane to detection performance as real world data collections have to encounter many of the same issues. This section will be devoted to identifying the issues related to data acquisition and the methods we used for each of our detectors. This is not meant to be a full comparison of all the different ways to process hyperspectral data. This comparison is only meant to help understand whether our hypothesis is valid. The following sections identify the experimental design and provide results for three experiments measuring endmember sensitivity, separation performance, and overall ROC performance. Table I summarizes these selections for each of the detectors (AMSD, ACE, HSD, and HUD) used in our experiments.

Table 9: Subpixel Experiment Details

Detector	Background Model	Background Signatures	Target Signatures	Abundance Constraints
AMSD	Structured	Eigenvectors of the Image Correlation Matrix	MODTRAN	No
ACE	Unstructured	Multivariate Normal with Global Covariance	MODTRAN	No
HSD	Structured	Iterative Error Analysis & Global Covariance	MODTRAN	Yes
HUD	Unstructured	Iterative Error Analysis & Global Covariance	MODTRAN	Yes

5.3.1. Experimental Design

These experiments require imagery, background signatures, target signatures, and ground truth information. The imagery used for these experiments comes from the Sensor X data described in Chapter 2. From this sensor, we used Images 1 through 6 because these images contain subpixel targets. The other images are full or multi-pixel targets which provide little challenge for the detectors.

As indicated in Chapter 4, we used two background endmember extraction techniques. The most significant eigenvectors of the global image correlation matrix were used as the “endmembers” for the AMSD algorithm as documented in [71]. Since these do not produce physically meaningful endmembers, we used the IEA algorithm for the hybrid detectors [77]. Additionally, we used the image covariance matrix for the hybrid detectors to whiten the data. In all cases, the endmembers and covariance matrices were estimated from the entire image. We also tried local estimates, but these provided results no better than using global estimates.

To choose the number of endmembers for each detector, we first extracted up to 150 endmembers for AMSD and 60 endmembers for the hybrid detectors. While we could have used our estimation techniques from Chapter 4, we decided to identify

the ideal cases for these results to present the best performance possible for each of the detectors. The concept of best performance turned out to be trickier than we first imagined. For images where we had targets, the best performance was defined as the number of endmembers that maximized the probability of detection while minimizing the number of false alarms. In cases where perfect separation was achieved between targets and false alarms, the best performance was defined using a minimax criterion where the clutter with the highest detection score was minimized. The same minimax criterion was applied to the cases where no targets were present. This method provided the best results independent of detector type both in terms of separation of targets and clutter and setting a fixed threshold for ROC curves.

The target information we received from NVESD were measured in units of reflectance. As discussed in Chapter 3, the images are measured in terms of radiance. There are three approaches to overcome this mismatch: use target signatures directly from the image for the experiments, convert the images to reflectance, or convert the targets to radiance.

Because the images only contained sub-pixel targets, we could not directly use target signatures from the image. If we did, the signatures would be corrupted with background and bias our results. Moreover, using target signatures from the image reduces the pool of targets. Those targets would have to be dropped from the analysis as any target pulled from the imagery would be guaranteed to be detected biasing the results.

Because of these reasons we turned to the atmospheric compensation techniques documented in Chapter 3. For the analysis in this chapter, we relied on the

model-based method MODTRAN to generate the target signatures. We did this to remove variability that may have been introduced using the ARRT method. This was especially true for Target 4 where the low reflectance signature made estimation using in-scene methods difficult.

Ground truth was used to create background and target objects. Following the procedures in Chapter 2, we applied a cluster threshold to each detector output to guarantee 1% of the pixels were above the threshold. This threshold was used knowing that the number of targets in the image was far less than 1% of the pixels in the image. Adjacent pixels above the threshold were assigned to the same cluster. In each cluster, the maximum detection score was assigned as the cluster detection score. These clusters' positions were then compared to ground truth information to label the clusters as either target or false alarms.

Fill factors for the experiment ranged from 10% to 60%. Fill factor describes the percent of the pixel that is occupied by target signature. Fill factors assume that the target lies exactly within the pixel. In numerous cases, subpixel targets can lie across pixel boundaries or be obscured by the competing environment (e.g. tall grass) generating fill factors in the image that are much smaller than expected.

5.3.2. Endmember Sensitivity Analysis

This experiment measures how sensitive the AMSD, HSD, and HUD algorithms are to the number of endmembers. In our experiments, we have ground truth information and hence can determine the “best” number of endmembers as defined in the previous section. In real world applications, this knowledge is not available to us; hence, we have to rely on algorithms to estimate the correct number of endmembers without the associated ground truth. As shown in Chapter 4, these

algorithms can have significant errors. Therefore, detectors that are insensitive to these estimation errors are highly desirable.

For this experiment, we measured the probability of detection and number of false alarms at varying numbers of endmembers from one to 60 across all images, targets, and detectors. We stopped at 60 endmembers because in all cases, the performance for all detectors on all targets degraded well before reaching this number and continued to degrade as will be shown in our results. The only exception to this rule was the number of endmembers used for AMSD. Additional experimentation showed that we needed to extract as many as 150 endmembers to provide good detection results.

We present the results in two ways. First, we provide an example to show how the false alarm density varies with the number of endmembers and type of detector. The results for this experiment are in Figure 25 which shows the performance of the AMSD, HSD, and HUD algorithms on Image 1 and Target 2. We chose this image and target type because it is indicative of the entire set of results we produced. The figure shows the number of false alarms for varying numbers of endmembers on each detector. We did not include the P_d figures simply because all detectors were able to achieve nearly 100% P_d across all numbers of endmembers. Therefore, good performance on this test is achieved if a minimal number of false alarms are detected across multiple numbers of endmembers. This indicates that the detector is partially insensitive to the number of endmembers chosen.

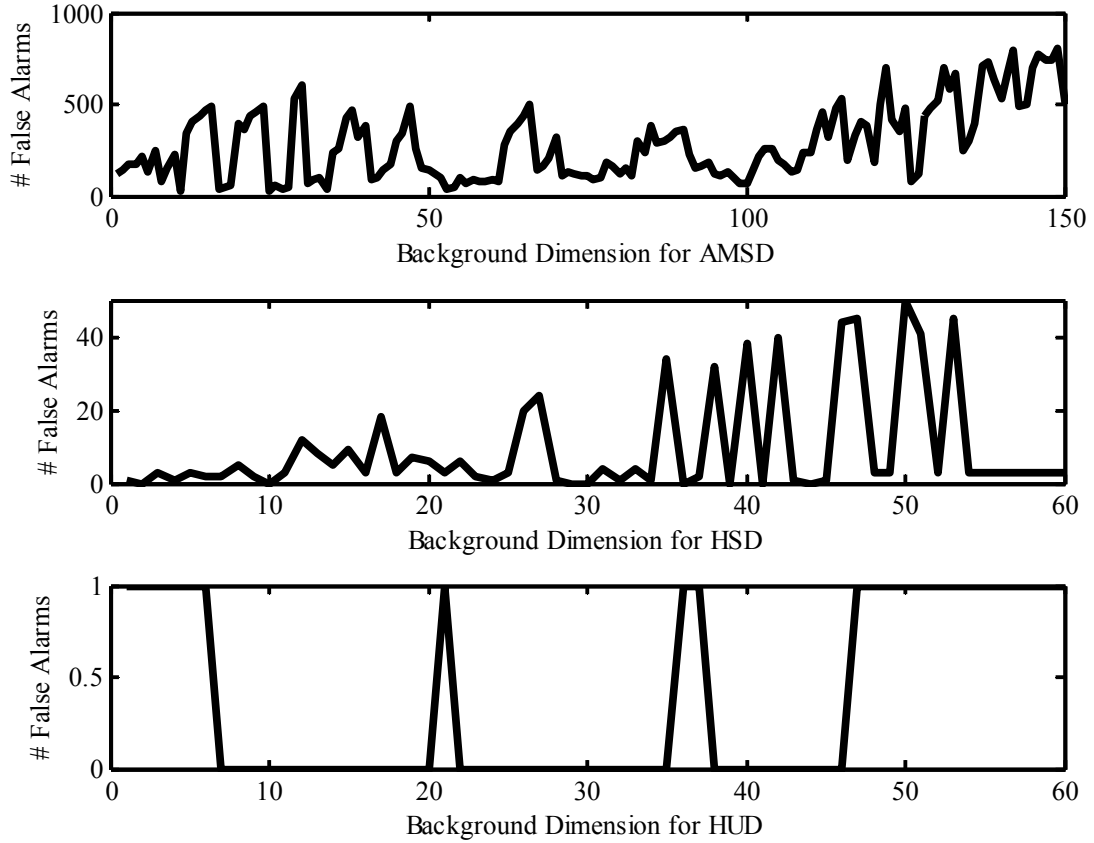


Figure 25: Graphical Comparison of Endmember Sensitivity

Figure 25 shows the hybrid algorithms are more insensitive to the number of endmembers than AMSD. The AMSD results are random and lack the general trend seen in the HSD and HUD results. When using AMSD, even slight changes in endmembers can produce dramatically different results varying from 27 false alarms to nearly 800. The HSD algorithm results show that endmembers numbering less than ten tend to produce better results. Also, the HSD results at higher number of endmembers do not vary as greatly as the AMSD figures. Instead, the worst case number of false alarms is limited to 50. HUD is the best in terms of being insensitive to the number of endmembers. This algorithm provides excellent performance regardless of the number of endmembers. The data shows that the hybrid detectors are

insensitive to the number of endmembers with HUD being nearly independent of them.

Figure 25 only shows the results for one image and one target type. To verify that this occurs for all target types and images, we put together Table 10 that contains the number of times the best performance was achieved across the 60 endmembers for each detector. The more insensitive a detector is to the number of endmembers, the higher the number. Best performance is defined as the instances that achieve 100% P_d with the lowest number of false alarms. Note that this could mean that the lowest number of false alarms is greater than zero. Results are only posted for images where the target is present.

Table 10: Endmember Sensitivity Results

Target	Image	AMSD	HSD	HUD
1	1	54	35	60
	4	39	34	60
2	1	1	8	37
	4	3	42	60
3	2	2	36	59
	3	21	42	60
	5	1	1	2
	6	1	1	59
4	2	0	0	0
	3	0	0	0
	5	0	0	0
	6	0	0	0

The results in Table 10 support the results from the first experiment. Target 1 is the easiest of the targets and this is demonstrated by the high numbers achieved with all the detectors. As the targets become more difficult to identify though, the results start to diverge. AMSD performance drops to single digits as target difficulty increases. HSD maintains good numbers until the hardest images where it too drops

to single digits. HUD fairs the best maintaining nearly perfect performance on all the images except a few. These experiments show the hybrid detectors are partially insensitive to the number of endmembers selected. Since the true number of endmembers is rarely if ever known, detectors with this insensitivity have a significant advantage over those that do not.

The only exception to the rule is Target 4 where none of the detectors are able to achieve 100% Pd. In this case, the performance is poor independent of the number of endmembers. The most likely cause is the target is so weak that target characterization methods are not correctly modeling the signature. This mismatch causes all detectors to perform poorly.

5.3.3. *Separability Analysis*

Having shown that the hybrid detectors are more insensitive to the number of endmembers selected, the question remains whether they provide improved detection performance over their AMSD and ACE counterparts. This set of experiments answers this question using figures that show the separability between target and background for each image and detector type. The figures were patterned after those found in [69]. These graphs are very useful because they can be used even when few targets are present. This allows us to measure the performance of the detectors on each image and target type.

The figures for each target type are shown in Figure 26 through Figure 29. Each figure contains four sub-figures. Each sub-figure contains black and gray vertical bars. The black bars show the range of detection values for the background. The gray bars show the range of detection values for the targets. Ideally, these bars should not overlap indicating the targets are completely separable from the

background. In cases where overlaps do occur, a number is posted above the black bar. This is the number of false alarms that occur if all of the targets are detected. Within any sub-figure, the ranges of the targets and background can be compared across images to see the consistency of the detector. A good detector will consistently suppress the background into a similar range of values while separating the targets.

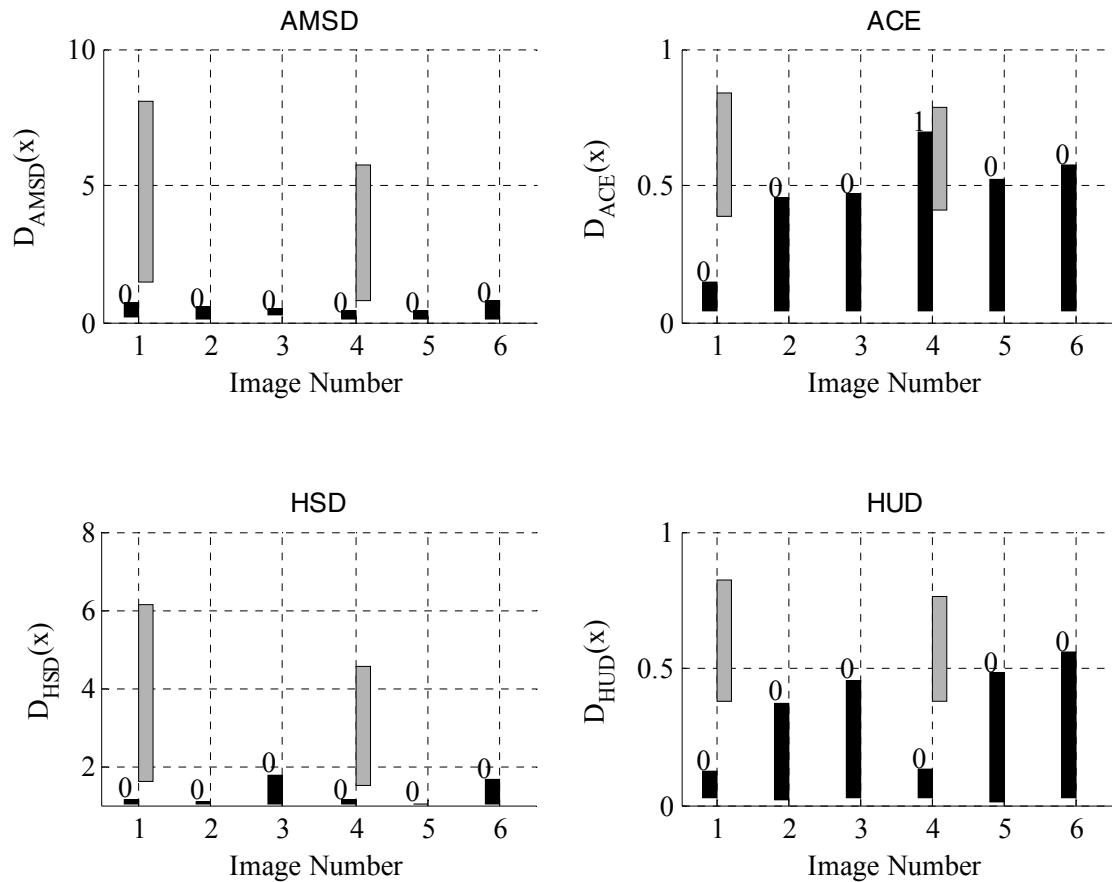


Figure 26: Separability Analysis for Target 1

Figure 26 shows the results for Target 1. This is the easiest target due to its white color that makes it very different from the surrounding background. All the detectors perform well with only ACE picking up one false alarm on Image 4. The structured detectors however perform better than their unstructured counterparts. The ACE and HUD algorithms do separate the target from the background, but have

difficulty suppressing the background in the images where the targets are not present. The structured detectors do not suffer from this problem and suppress the background nearly equally across all images. AMSD has a slight advantage over HSD on Images 3 and 6 where the background values have been compressed a bit farther than with HSD. Nevertheless, all detectors show good performance on this target.

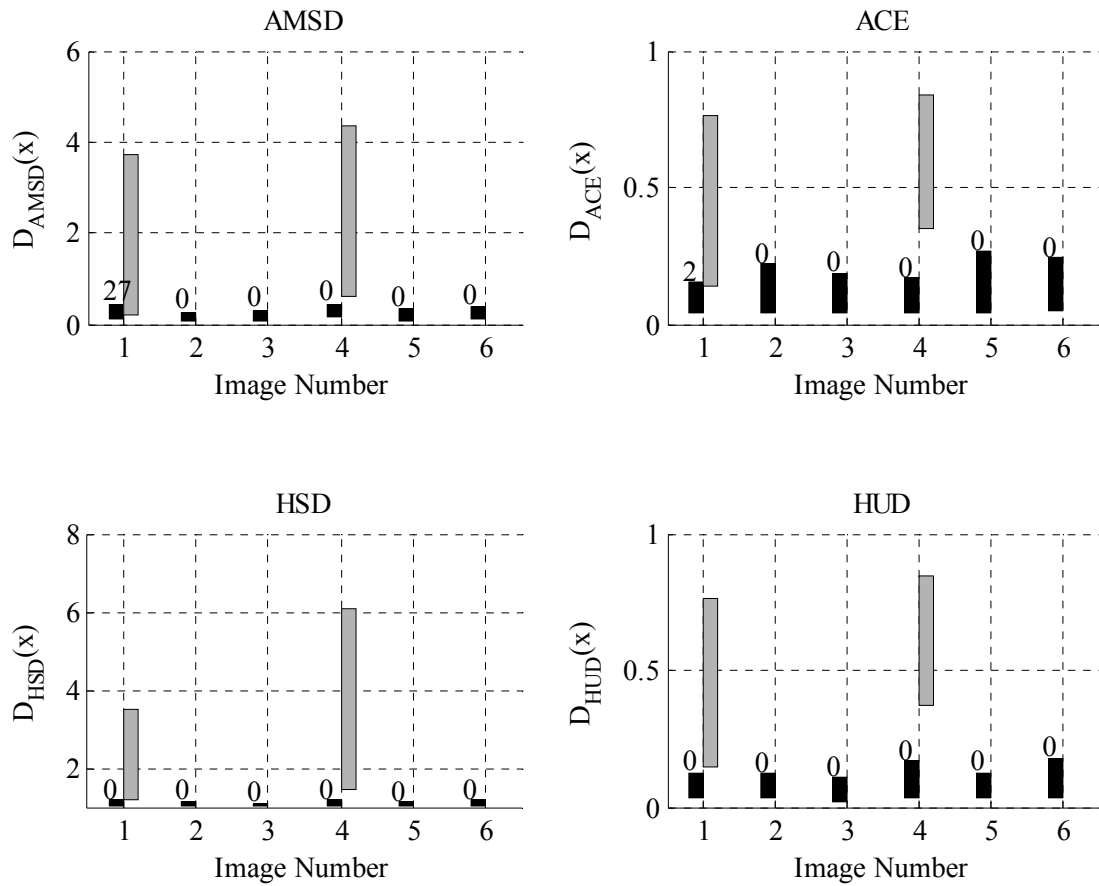


Figure 27: Separability Analysis for Target 2

Figure 27 shows the results for Target 2. This target is painted green and although larger than Target 1 is harder to separate from the green background. It is with this target that the hybrid detectors begin to show a slight performance advantage over the standard detectors. The hybrid detectors maintain zero false alarms across all images as was the case with Target 1. AMSD however picks up 27

false alarms on the first image and ACE picks up 2 false alarms on the same image. The hybrid detectors also do a better job of suppressing the background into similar ranges across the images. AMSD does as well but has the aforementioned 27 false alarms. ACE is the only detector where the background values vary significantly across the images.

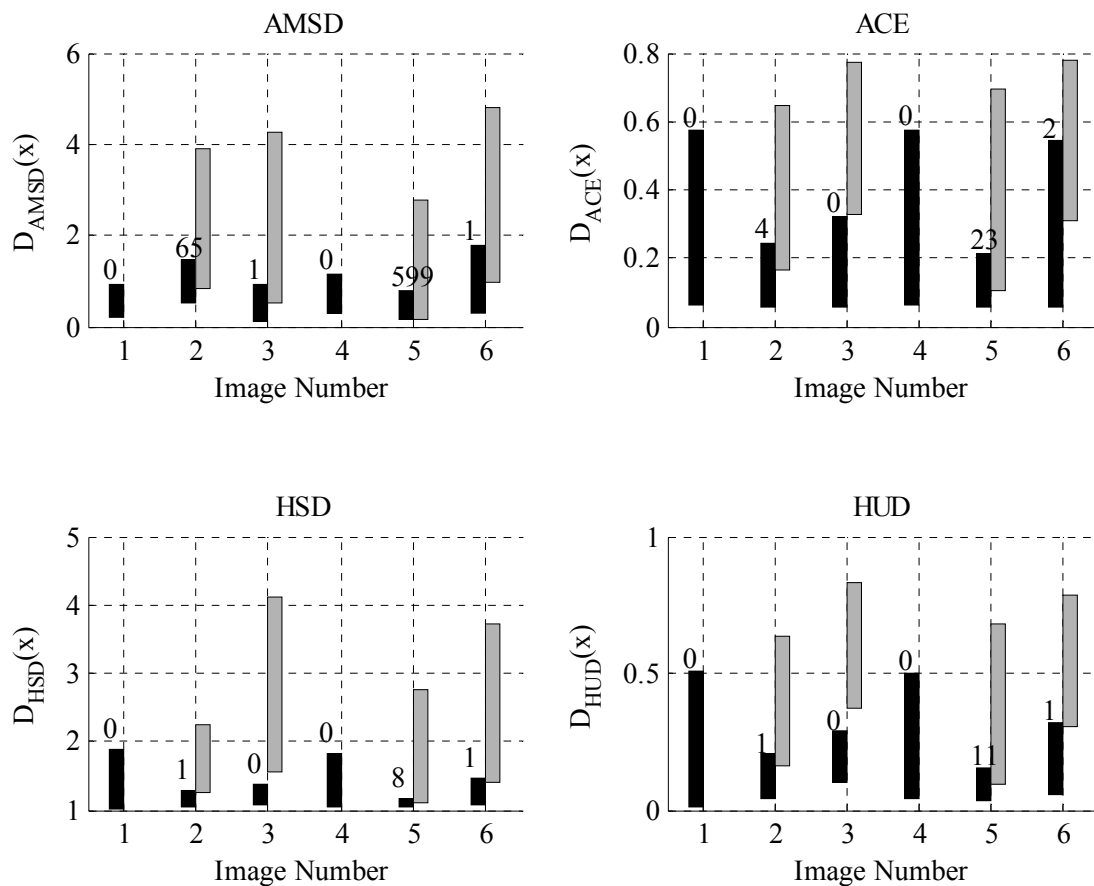


Figure 28: Separability Analysis for Target 3

Figure 28 shows the results for Target 3. This target has multiple reflectance signatures which indicate a significant variability of the spectral signature. Because of this variability, all the detectors have difficulties with this target. The background is no longer being compressed to the same range of values for any detector although the structured detectors do fair better than their unstructured counterparts. The key is the

number of false alarms. AMSD achieves 666 false alarms across all images. ACE drops this number to 29. HUD further reduces the number to 13 while HSD performs the best with only 10 false alarms. These numbers are remarkable in that the hybrid detectors have provided 66 times less false alarms than AMSD and 3 times less false alarms than ACE. When one considers that the hybrid detectors are also the most insensitive to the number of endmembers selected, the performance gains become much more significant.

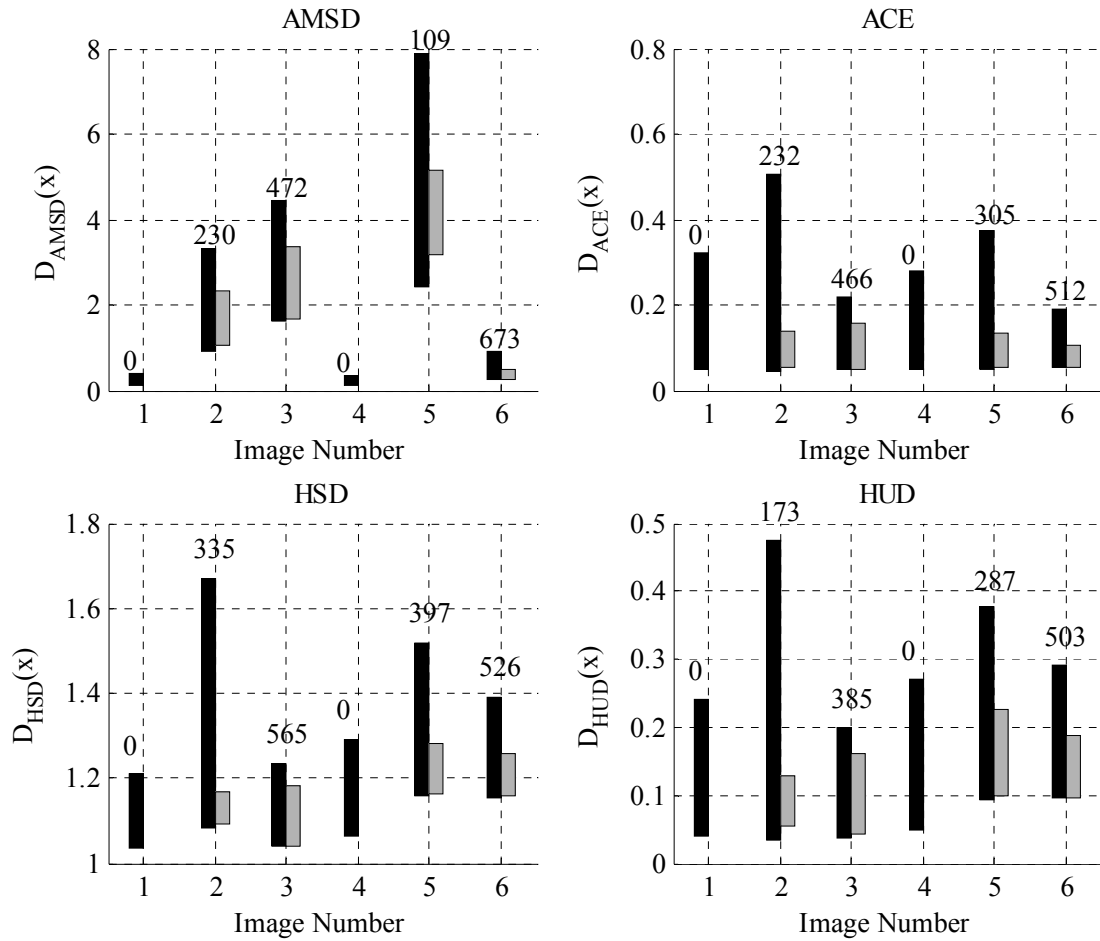


Figure 29: Separability Analysis for Target 4

Figure 29 shows the results for Target 4. As expected, all of the detectors have difficulty with this weak target. This is the only target where the hybrid detectors

show no improvement over the statistical detectors AMSD and ACE. The most likely cause of this result is the incorrect modeling of the target radiance signature as noted in Chapter 3. When the estimated target does not match the target signature in the image, no signature based detector is going to perform well. None of the detectors are able to detect 100% of Target 4 in any of the images. Therefore, this target is not a good example for comparing the different subpixel target detectors, but it does support the need for good target characterization.

5.3.4. Receiver Operating Characteristics

In our separability analysis, we argued that some detectors did a better job consistently pushing the background values into a similar region across all the images. A good way to measure this consistency is to use a ROC curve. The ROC curves we generate are for a single detector and single target across all images. This provides enough target returns to make each ROC statistically significant. Note that a ROC measures the average performance for a fixed threshold across all images; therefore, detectors that consistently separate the targets and background into similar detection values across each image will perform better than those that do not. Theoretically, the CFAR algorithms AMSD and ACE should provide such performance. Our interest is whether the hybrid algorithms will meet or exceed the results of the CFAR algorithms thus giving them CFAR-like properties although this fact cannot be proved theoretically.

Figure 30 shows the ROC curves for Target 1. As expected from our separability analysis, the structured detectors outperform the unstructured detectors. AMSD does have a slight performance improvement over HSD, but the results show

the hybrid detectors are achieving the same CFAR performance as the standard detectors.

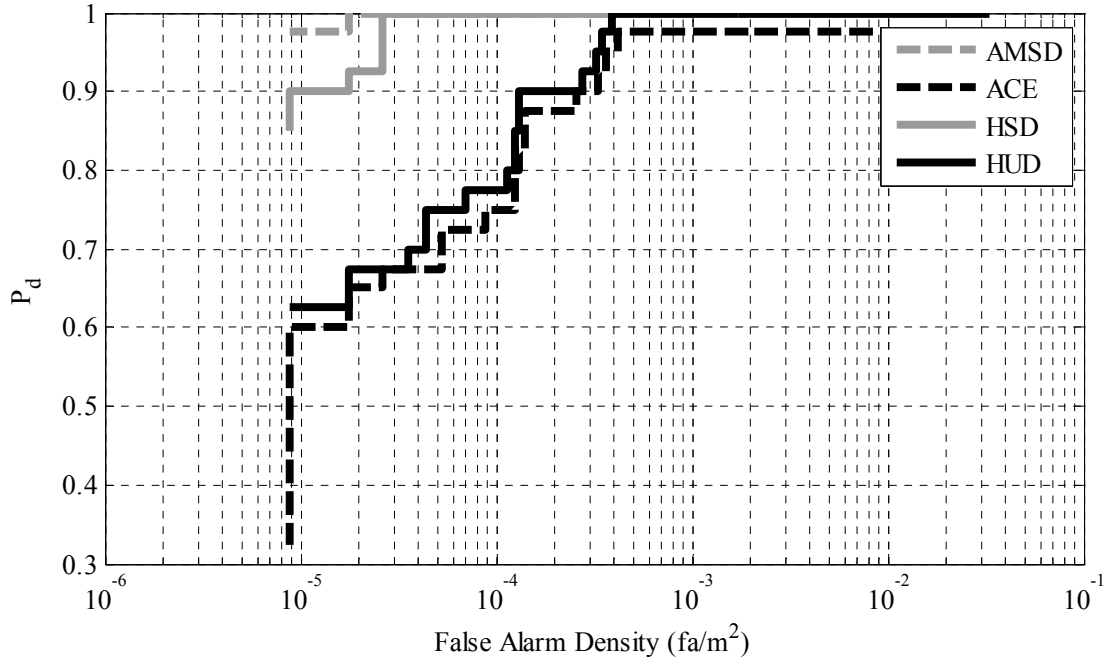


Figure 30: Subpixel Detection ROC Curves for Target 1

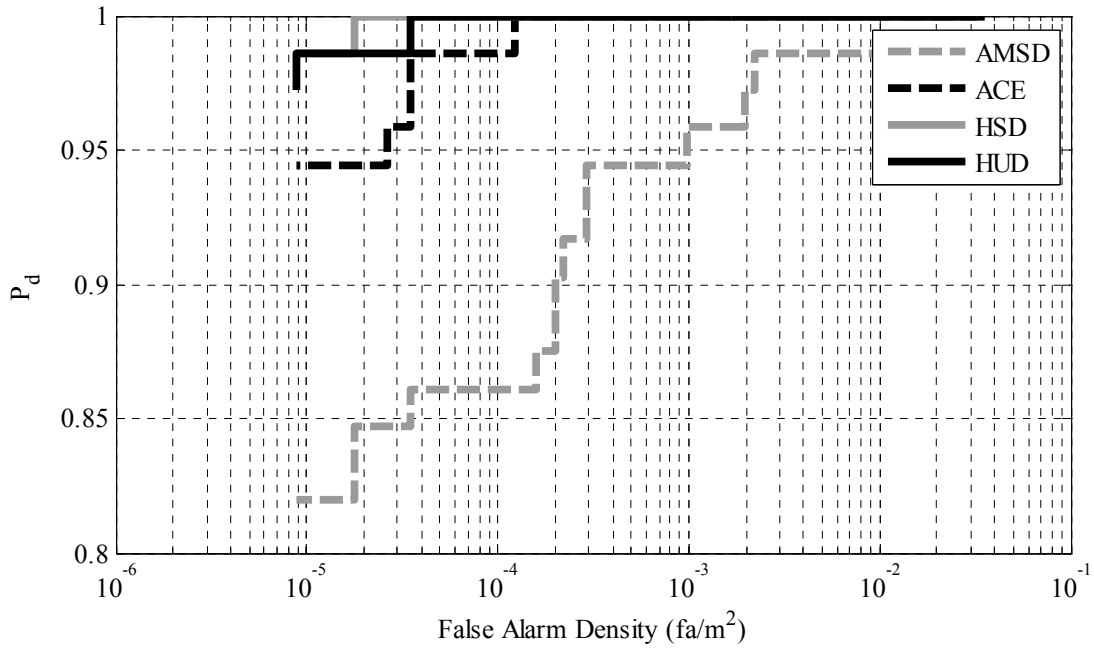


Figure 31: Subpixel Detection ROC Curves for Target 2

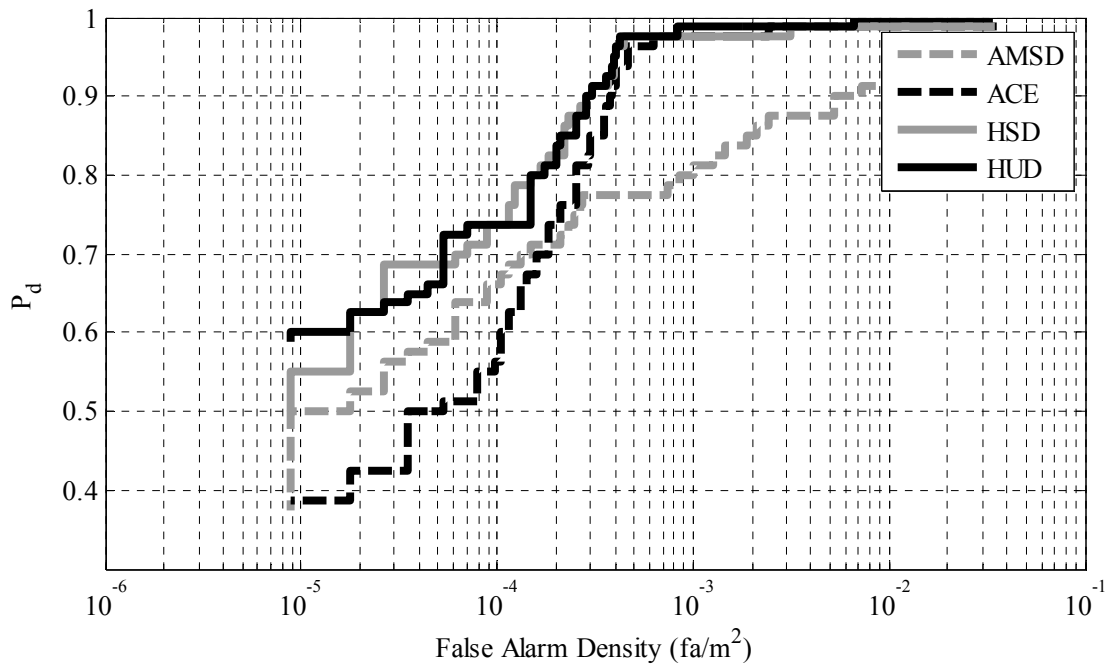


Figure 32: Subpixel Detection ROC Curves for Target 3

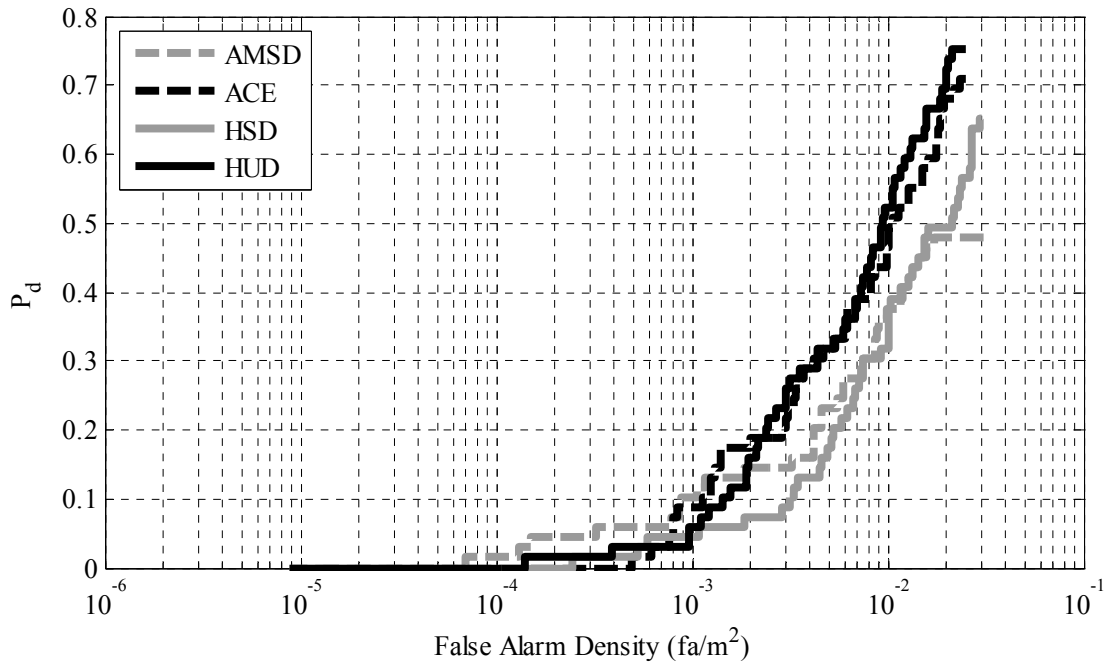


Figure 33: Subpixel Detection ROC Curves for Target 4

Figure 31 shows the ROC curves for Target 2. The hybrid detectors are slightly better than their standard counterparts. While the figure seems to show a great

improvement in performance, the range of P_d values is measured from 0.8 to 1.0. This gives the impression of much better performance. Nevertheless, the hybrid algorithms are again performing as well if not better than their CFAR counterparts.

Figure 32 shows the ROC curves for Target 3. In the separability analysis, the hybrid algorithms showed great performance improvements over AMSD and ACE. What was noted in that section was that none of the detectors were able to suppress the background into a consistent range of values. The ROC curves show this fact. The hybrid algorithms are performing better than their CFAR counterparts, but the performance improvement is not as significant as in the separability analysis. The conclusion that can be drawn from this result is that the background and target are similar making the background harder to suppress. Nevertheless, the hybrid detectors are modeling the background better than AMSD and ACE which provides the gains in performance.

Figure 33 shows the ROC curves for Target 4. As expected, none of the detectors perform well. This is the only target for which the acceptable performance criteria of 50% P_d at 10^{-3} false alarms/m² is not met. As mentioned before, the reason is due to incorrect modeling of the target radiance signature.

5.3.5. Conclusions

Our set of experiments demonstrates the usefulness of the hybrid detectors. These detectors have a three-fold gain over their standard counterparts. First, they are tolerant of slight errors in the number of endmembers. Second, they show greater separability between targets and background – especially as the target becomes more difficult to detect. Third, they maintain a slightly more consistent threshold across the images than the known CFAR detectors AMSD and ACE. This result argues the

hybrid detectors' ability to better model the background and therefore detect subpixel targets.

What has not been mentioned so far is the efficiency of the hybrid algorithms. The algorithms require very little extra processing time when compared to either AMSD or ACE. ACE was perhaps the fastest of the detectors since we estimated the covariance matrix from the entire image. Results were also generated for ACE using local neighborhoods, but the performance showed little to no improvement over using the entire image. AMSD was nearly as quick as ACE except for the extraction of endmembers using an eigenvalue decomposition of the image correlation matrix. The hybrid detectors took the longest, but only because of the IEA endmember extraction algorithm. Once the endmembers were extracted, the performance was no different than that achieved with AMSD. The reason for this is the efficient FCLS algorithm which only took ten minutes to process an image when using 60 endmembers. With endmembers less than 20, the FCLS algorithm took less than a minute. Since most of the hybrid detectors prefer endmembers numbering less than 20, the processing times were similar to AMSD.

One final note is on the difference between the HSD and the HUD algorithms. Both of these algorithms performed well, but the HSD algorithm has a slight performance advantage. On all targets it was able to achieve false alarm densities smaller than HUD. HSD was also more consistent in suppressing the background into a similar range of detection values. The tradeoff is that the HSD algorithm is more sensitive to the number of endmembers. For example, the HSD algorithm requires an estimate of the number of endmembers that is close to the ideal. HUD on the other

hand can simply set the number of endmembers to some fixed number and achieve the same results for nearly all images. Therefore, the HUD algorithm is not as dependent on the number of endmembers, but has slightly lower performance than HSD because of this fact.

5.4. *Summary*

In this chapter we argue that better characterization of the background through physics-based knowledge can improve subpixel detection performance. To this end, we develop two hybrid detectors which use physically meaningful endmembers and abundances within a statistical hypothesis test. We compare these detectors to their purely statistical counterparts AMSD and ACE.

Our results show that the improved background models of the hybrid detectors provide improved performance in three different ways. First, the hybrid detectors are less sensitive to the number of endmembers used. Thus, endmember estimation algorithms can allow some error without significantly degrading subpixel detection performance. Second, the hybrid algorithms provide better separation between the targets and background per individual image. This is especially the case with weaker targets like Target 3 where AMSD and ACE have false alarm densities well over 30 compared to 10 for the hybrid detectors. Finally, the hybrid detectors provide a more consistent separation of target and background that leads to improved ROC performance.

While this research shows the importance of modeling the background on subpixel target detection algorithms, further research is required. On Target 3, the hybrid detectors did outperform their statistical counterparts, but Figure 28 shows that the background detection scores can still vary significantly from image to image. One

way to counteract this phenomenon is to better characterize the background using more appropriate density functions or non-parametric techniques in conjunction with physics-based knowledge. Another means to counteract this phenomenon is to use adaptive threshold techniques. Either way, our research suggests much more can be done to model and understand the complex background inherent in hyperspectral imagery to improve subpixel target detection performance.

Chapter 6: Adaptive Detection Thresholds via Extreme Value Theory

Subpixel detectors present a significant challenge in determining the detector threshold for a desired probability of false alarm. For example, the most common threshold estimation method is a theoretical calculation for used for CFAR detectors. CFAR detectors are designed such that the distribution of the detector given the background is independent of any estimates needed to derive the detector [70]; therefore, the conditional background distribution is independent of the data. This independence of the clutter distribution from the data allows a theoretical calculation of a fixed false alarm density α_0 . CFAR detectors achieve this goal by making an assumption about the underlying distribution of the data. Typically this assumption is that the underlying distribution is a normal distribution (or at least any zero-mean elliptically contoured distribution [57]), which makes the mathematics tractable enough to determine the detector's statistical distribution. Additionally, CFAR detectors typically assume independent and identically distributed (iid) samples. For instance, a standard detector for HSI data is the Adaptive Cosine Estimate (ACE) detector which assumes the underlying distribution is multivariate normal [58]. ACE is a CFAR detector whose threshold can be calculated theoretically for a desired false alarm density. In practice though, HSI data has been shown to be rarely multivariate normal [103] and hence any theoretically calculated threshold for the ACE detector is most likely inaccurate.

In recent publications, the use of elliptically contoured distributions has been explored to model the outputs of detectors [69]. This method is similar to the theoretical threshold calculations for CFAR detectors except the method models the

output of the detector as an elliptically contoured distribution. The detector data is then used to estimate parameters which in turn provide a distribution from which a detection threshold can be theoretically calculated. The usefulness of this method is currently being investigated, but its applications are limited to CFAR detectors. This prevents us from using these techniques for our hybrid detectors where the output distribution is difficult at best to determine due to the non-negativity constraints. Therefore, we must rely on methods that directly use the output detection statistics.

A standard non-parametric approach for determining the desired detector threshold is to use order statistics. The detector output is sorted in descending order to create an ordered list. The number of detection values N is multiplied by the desired α_0 and rounded to the nearest integer. This integer is used to identify the position in the ordered list that will be used as the detection threshold. The strength of this approach is that any detector output can be used – not just those that are CFAR. Even if the detection threshold varies significantly from image to image, the use of this method adjusts the threshold automatically to track such deviations. Unfortunately, the method is very sensitive to outliers when low false alarm densities are required. For example, a typical detection image will contain both targets and clutter. The order statistic algorithm will count the targets as clutter and this will skew the detection threshold. We can think of this as a Monte Carlo (MC) method where instead of estimating the probability of false alarm density from the detector samples, we use the samples to estimate the threshold for a desired false alarm density. In subsequent discussions, we will call this the MC method.

Another method of determining the detector threshold is based on importance sampling (IS). Importance sampling is a forced Monte Carlo method that is used to simulate rare events [101]. IS has been mostly used to test system responses to rare events in an efficient manner. There are a number of papers that prove its ability to provide unbiased estimates of rare event probabilities with low variance [91][99][102]. These rare events simulate the distribution tails of the system and hence are closely related to the design and measurement of detectors.

Srinivasan showed that IS could be used to determine a detector threshold for a desired fixed false alarm probability α_0 [101]. This method is called inverse importance sampling. Initially, these thresholds were determined for standard background distributions that a detector may encounter such as the normal, Rayleigh, or Weibull distributions. Bucklew extended this research to handle situations where the underlying probability density function was unknown [17]. Unfortunately, these methods are designed for sums of random variables. In [101], Srinivasan shows that blind importance sampling when applied to data from a single random variable provides no gains over MC methods. Since the detector output is from a single random variable, blind IS methods are not ideal.

Therefore, we turn to the use of Extreme Value Theory (EVT). EVT concerns problems where the probability of a rare event must be estimated even if such a rare event has never occurred [39]. This type of research has wide applicability in such fields as climatology [100], detection theory [74], anomaly detection [89], and financial analysis [25]. It is in the last field where most of the theory has been applied to estimate stock market anomalies, insurance rates for catastrophic events, and

management of risk. These applications are very similar to our problem of estimating a threshold for rare events even if they have not occurred. This makes EVT a variance reduction technique similar to IS, but far more applicable to wider class of problems [38].

In target detection, the presence of targets can significantly impact the performance of threshold estimates. A variety of methods have been developed to remove outliers (e.g., isolation of target returns from the background) [47]. These methods vary widely from simple sample statistics to advanced classification techniques based on Support Vector Domain Descriptions [108]. Interestingly, EVT theory can also be used to identify outliers in a data sample [89]. Thus, EVT can both estimate detection thresholds for a given false alarm density and simultaneously be used to remove the influence of outliers on the sample.

Therefore, we present a novel adaptive threshold technique based on extreme value theory. The new technique is able to set thresholds for desired probabilities of false alarm densities similar to the MC technique. Unlike the MC technique, we develop an outlier rejection capability using the Generalized Pareto Distribution (GPD) that can identify samples that do not belong to the same distribution as the background. These outlier samples can be removed such that desired false alarm densities in the presence of target returns can be calculated with some confidence. The rest of the chapter is structured as follows. Section 6.1 presents an overview of Extreme Value Theory. Section 6.2 describes our adaptive threshold algorithm based on GPD estimates. Experimental results are given in Section 6.3. A summary concludes the chapter in Section 6.4

6.1. Extreme Value Theory

6.1.1. *The Fisher-Tippett Theorem*

Assume there is a set $X = \{x_1, x_2, \dots, x_m\}$ of m i.i.d. samples drawn from the same unknown and continuous cdf $F(x)$. Denote the maximum of the set X as

$$x_{(m)} = \max(X). \quad (70)$$

with cdf

$$H(x) = [F(x)]^m. \quad (71)$$

Fisher and Tippett [28] show that if $H(x)$ is stable in the limit as $m \rightarrow \infty$, then an affine transformation exists such that

$$x_{(m)} \stackrel{d}{=} \sigma_m x + \mu_m \quad (72)$$

for a given scale parameter σ_m and location parameter μ_m . Equation (72) states that the maximum of the set X converges in distribution to the affine transform. Using the affine transformation given, Fisher and Tippett show that

$$H(x_{(m)} \leq x) \stackrel{d}{=} H(\sigma_m^{-1}(x - \mu_m)) \quad (73)$$

the normalized form is the only form for the limit distribution of $x_{(m)}$ given any $F(x)$.

Now assume that $H(x)$ is a non-degenerate limit distribution for normalized maxima of the form $\sigma_m^{-1}(x - \mu_m)$, then $H(x)$ is only one of three forms. This theorem is the famous Fisher-Tippett theorem [28] and is the foundation for extreme value theory. Denoting $y_m = \sigma_m^{-1}(x - \mu_m)$, the “reduced variate”, the three forms are

$$\begin{aligned}
H_1(y_m) &= \exp(-\exp(-y_m)) \\
H_2(y_m) &= \begin{cases} 0 & \text{if } y_m \leq 0 \\ \exp(-y_m^{-\alpha}) & \text{if } y_m > 0 \end{cases} \\
H_3(y_m) &= \begin{cases} \exp(-(-y_m)^\alpha) & \text{if } y_m \leq 0 \\ 1 & \text{if } y_m > 0 \end{cases}
\end{aligned} \tag{74}$$

for $\alpha > 0$ which are the Gumbel, Fréchet, and Weibull distributions respectively. What this theorem states is that as $m \rightarrow \infty$, the maximal distribution $H(x)$ is in the domain of attraction of one of the three limit forms in (74) for any $F(x)$. Therefore, much like the central limit theorem for sums of random variables, the Fisher-Tippett theorem provides a known limiting distribution for the maxima from any set of i.i.d. samples.

6.1.2. EVT for the Exponential Class

Most research has focused on the type I or Gumbel distribution. This limiting distribution occurs for all samples that are drawn from a distribution in the exponential class [35][39] which contains such well-known distributions as the normal, lognormal, and K distributions. A number of researchers have developed theory to identify whether data samples belong in the exponential class such as Gumbel [39], Gnedenko [35], and von Mises [112]. From this theory, Weinstein [114] introduced the generalized extreme value theory (GEVT) such that

$$\lim_{m \rightarrow \infty} H\left((a_m^\nu + c_m y)^{1/\nu}\right) = \exp(-\exp(y)) \tag{75}$$

where $a_m > 0$, $\nu > 0$, and

$$x_m = (a_m^\nu + c_m y)^{1/\nu}. \tag{76}$$

When considering tail estimates based on data from the exponential class, the Gnedenko criterion states that (75) holds if and only if

$$\lim_{n \rightarrow \infty} n \left\{ 1 - F \left((a_n^v + c_n y)^{1/v} \right) \right\} = -\exp(y), \quad \forall y. \quad (77)$$

Using (75) through (77), we can estimate the tail of the unknown exponential class $F(x)$ by

$$Q(x) = 1 - F(x) = \frac{1}{n} \exp \left(\frac{-(x^v - a_n^v)}{c_n} \right). \quad (78)$$

Having defined the unknown tail probability, we need to estimate the four parameters: a_n , c_n , v , and n . Guida, Iovino, and Longo present a way to find these parameters using numerical optimization of the maximum likelihood estimates [38].

These estimates are

$$\hat{c}_n = \frac{1}{L} \sum_{i=1}^L x_{n(i)}^{\hat{v}} - \frac{\sum_{i=1}^L x_{n(i)}^{\hat{v}} \exp(-x_{n(i)}^{\hat{v}}/\hat{c}_n)}{\sum_{i=1}^L \exp(-x_{n(i)}^{\hat{v}}/\hat{c}_n)}, \quad (79)$$

$$\hat{a}_n = -\hat{c}_n \ln \left(\frac{1}{L} \sum_{i=1}^L \exp(-x_{n(i)}^{\hat{v}}/\hat{c}_n) \right), \quad (80)$$

and

$$\hat{v} = L \left\{ \sum_{i=1}^L \left\{ \ln x_{n(i)} + \left(\frac{x_{n(i)}^{\hat{v}} \ln x_{n(i)}}{\hat{c}_n} \right) \left(\exp \left(\frac{x_{n(i)}^{\hat{v}} - \hat{a}_n}{\hat{c}_n} \right) - 1 \right) \right\} \right\}^{-1} \quad (81)$$

where $x_{n(i)}$ is the maximum value from the i^{th} set of n samples. These can be iteratively solved using numerical techniques such as the Kimball procedure [39].

The only other parameter to be estimated is n . Unfortunately, this parameter cannot be estimated using MLEs. Instead, Guida, Iovino, and Longo perform a number of trials to see the effect of this parameter on the final solution [38]. Their results show that n should be on the order of tens of samples to maximize the number

of L sets. If n becomes to large, L decreases leading to poor estimates of the tail distribution.

6.1.3. Generalized Pareto Distribution

Pickands [79] noted that classical EVT (Fisher-Tippett theorem) has a number of difficulties when applied in practice. First, most research has focused on only one of the three limiting distributions – namely, the distribution for data from the exponential class as noted in the previous section. Unfortunately, if the data does not come from the exponential class, a practitioner must use his/her intuition and subjective reasoning to choose the correct parametric model. Second, classical EVT requires partitioning the data into n set of m samples. As noted in [38], there is no direct way to identify the best partitioning a-priori. To this end, Pickands [79] and Balkema and de Haan [5] introduce a new way to estimate the tail of a distribution based on modeling the distribution of samples above some high threshold.

Following the work of Pickands [79], assume that we have n i.i.d. samples from a continuous and unknown distribution $F(x)$. Pickands assumes for some c , $-\infty < c < \infty$, there exists

$$\lim_{n \rightarrow x_\infty} \inf_{0 < a < \infty} \sup_{0 \leq x < \infty} \left| \frac{1 - F(u + x)}{1 - F(u)} - \exp \left(- \int_0^{x/a} [(1 + ct)_+]^{-1} dt \right) \right| = 0 \quad (82)$$

where $x_\infty = \text{greatest lower bound } \{x: F(x) = 1\} = \text{lowest upper bound } \{x: F(x) < 1\}$, and $y_+ = \max(0, y)$. For any u and x , the $[1 - F(u+x)]/[1 - F(u)]$ is the conditional probability that an observation is greater than $x+u$ where u is some high threshold. Therefore,

$$P(X | X \geq u) = 1 - G(x) = \exp \left(- \int_0^{\frac{x/a}{(1+ct)_+}} [(1+ct)_+]^{-1} dt \right). \quad (83)$$

Von Mises [112] showed for EVT that the extremal distribution functions have the form

$$\Lambda(x) = \exp \left(- \int_0^y [(1+ct)_+]^{-1} dt \right), \quad y = \sigma_m^{-1}(x - \mu_m). \quad (84)$$

Therefore, $P(X \geq u)$ is in the domain of attraction of the classical EVT distributions without having to partition the data into n sets of m samples.

If $F(x)$ is continuous, then $G(x)$ is a generalized Pareto distribution (GPD) of the form

$$G(x) = \begin{cases} 1 - \left(1 + c \frac{x}{a} \right)^{-\frac{1}{c}} & \text{if } c \neq 0 \\ 1 - e^{-\frac{x}{a}} & \text{if } c = 0 \end{cases} \quad (85)$$

for all x such that $0 < x < \infty$. Depending on the shape factor c , the GPD embeds a number of other distributions. When $c = 0$, the GPD is an exponential distribution. When $c > 0$, the GPD is the ordinary Pareto distribution. When $c < 0$, the GPD is the Pareto II distribution. Pickands also shows that the estimated GPD is consistent and converges in probability to the true tail distribution such that

$$\lim_{n \rightarrow \infty} P \left\{ \sup_{0 \leq x < \infty} \left| \frac{1 - F(\hat{u} + x)}{1 - F(\hat{u})} - [1 - \hat{G}(x)] \right| > \varepsilon \right\} = 0, \quad \forall \varepsilon > 0. \quad (86)$$

Therefore, the GPD is a consistent estimate of the tail distribution based on samples above some high threshold u for an unknown $F(x)$. The importance of this research is

that it removes the subjective selection of one of the extremal distributions in (74) and removes the need to partition the data set into n set of m samples.

6.2. EVT Adaptive Threshold Algorithm

Having described the main theorems for extreme value theory, we now proceed to describe how this theory can be used to estimate detection thresholds. Detection thresholds are typically set by fixing the threshold at a desired probability of false alarm (α_0). In CFAR detectors, this threshold can be calculated directly assuming the data fits the statistical distribution of the detector. In subpixel detection, the HSI data rarely fits the standard CFAR assumption of normal statistics. MC methods shown in (88) can be used to estimate the threshold from the data, but they are inaccurate for very small α_0 and are sensitive to outliers.

We can use the theory based on GPD to calculate the threshold for a tail distribution. Following the derivations in [33], we can redefine the unknown cdf as

$$F(x) = (1 - \Pr(X \leq t))F_t(x - t) + \Pr(X \leq t). \quad (87)$$

where t is a sufficiently high threshold. The probability that the set of data is less than t is easy to find using MC methods. The estimate is

$$\Pr(X \leq t) = \frac{N - n}{N} \quad (88)$$

where N is the total number of samples and n is the number of samples above t . Thus, the threshold needs to be high enough such that the remaining samples are in the tail of the distribution, but not so high that very few samples exist above the threshold. A good rule of thumb is to use either a threshold that captures 90% or 95% of the data. Note that this metric is a simple MC method and will provide unbiased, consistent estimates as the number of samples increase.

The remaining term in (87) is the cdf of the tail of the distribution $F_t(x-t)$. For this estimate, we use the GPD given in (85). To use the GPD, we must estimate the parameters a and c from the data. To perform this estimation, we calculate the log likelihood function from (85). To begin, we first calculate the probability density function (pdf) as the derivative with respect to x of (85) to obtain

$$g(x) = \begin{cases} a^{-1} \left(1 + c \frac{x}{a}\right)^{-1/c} & \text{if } c \neq 0 \\ a^{-1} \exp\left(-\frac{x}{a}\right) & \text{if } c = 0 \end{cases} \quad (89)$$

If we assume i.i.d. samples from the distribution, the likelihood equation is

$$g(X) = \prod_{i=1}^n g(x_i) \quad (90)$$

Taking the natural logarithm of (83), we obtain the log likelihood function

$$\log g(X) = \begin{cases} -n \log a - \frac{1+c}{c} \sum_{i=1}^n \log\left(1 + c \frac{x_i}{a}\right) & \text{if } c \neq 0 \\ -n \log a - \frac{1}{a} \sum_{i=1}^n x_i & \text{if } c = 0 \end{cases} \quad (91)$$

Unfortunately, the log likelihood equation is nonlinear and solving for each of the parameters results in coupled nonlinear equations. Therefore instead of trying to directly estimate the parameters using MLEs, we turn to the Nelder-Mead Simplex Method which is an implementation of unconstrained nonlinear optimization [62]. This method finds the minimum of a function; thus, instead of maximizing $\log g(X)$, we minimize $-\log g(X)$. Using this technique, we obtain estimates of a and c .

Having calculated all the parameters, we can rewrite (87) for the tail samples such that

$$\hat{F}(x) = 1 - \frac{n}{N} \left(1 + \hat{c} \frac{x-t}{\hat{a}} \right)^{-1/\hat{c}}. \quad (92)$$

Conversely, we can rewrite (92) to find the threshold for a given cdf value to obtain

$$t_\alpha = t + \frac{\hat{a}}{\hat{c}} \left(\left(\frac{N\alpha_0}{n} \right)^{-\hat{c}} - 1 \right) \quad (93)$$

where t_α is the threshold for a desired $\alpha_0 = 1-F(x)$ beyond threshold t .

This is a very useful result for our application. After setting a clustering threshold t , we can estimate a detection threshold t_α from the data samples for a desired α_0 value. The problem here as with the MC method is the GPD method assumes that all the data samples come from the same underlying distribution. In the case where targets are present, this assumption is invalid and suffers from the same problems as MC techniques.

The GPD method, however, is based on the knowledge that the tails of a distribution will converge in probability to the generalized Pareto distribution [79]. This only occurs though if the data samples come from the same distribution. When the data contains samples from multiple distributions, the tail will not converge to a GPD. We can use this knowledge to identify when target samples are present in the data and remove them before estimating a threshold for a desired α_0 .

To identify the presence of samples from two different distributions, we use the confidence bounds of the GPD. The idea is based on the fact that if the data comes from a single distribution, it should fall within the confidence bounds. Therefore, if we set 90% confidence bounds, 90% of the samples should fall between the bounds. If a higher percentage of samples fall outside these bounds, we hypothesize that the samples must come from multiple distributions.

To generate the confidence bounds, we rely on either numerical optimization or Monte Carlo simulation. Both provide reliable estimates of the GPD bounds, but we found the Monte Carlo simulations to be much quicker. To create these Monte Carlo estimates of the confidence bounds, we generate hundreds of random samples for each data sample of the GPD using the estimates found from (84). This provides a range of estimated $F(x)$ values at each data sample. The estimated samples are ordered. The confidence bound for the particular data sample is then calculated by taking the two estimated samples such that 90% of the remaining samples fall between them. This is done at every data sample to calculate the confidence bounds.

To help describe how we use the confidence bounds, we construct two simple examples. For the first example we generate 10,000 samples from a standard normal distribution. For the second example, we generate 9,900 samples from a standard normal distribution and 100 samples from a normal distribution with a mean value of 6 as “target” detections. We fit a GPD to the top 10% of the data for both examples. From these points, we estimate the tail cdf according to (83). We compare the results to the cdf calculated using MC techniques in (88) (also called the Kaplan-Meier empirical cdf [22]).

Figure 34 shows the estimated GPD with associated 90% confidence bounds compared to the empirical cdf for the first example. The solid gray red curves represent the 90% confidence bounds. The black points are the empirical cdf and the dashed gray line is the best fit using the GPD. The empirical cdf fits well between the confidence bounds having only 4 samples fall outside the bounds. This represents 0.4% of the samples which is much less than the 10% limits enforced by the bounds.

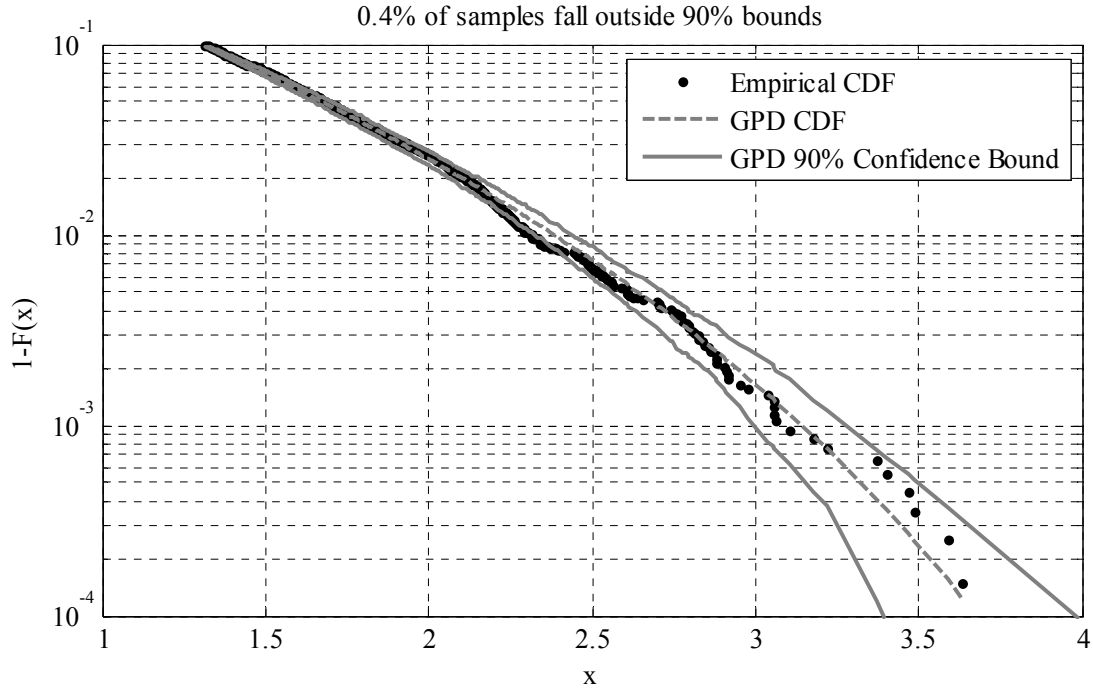


Figure 34: Comparison of the GPD to the Empirical CDF for Example 1

Figure 35 shows the estimated GPD with associated 90% confidence bounds compared to the empirical CDF for the second example. The empirical cdf falls well outside the bounds with over 30% of its samples beyond the 90% confidence limits. This example is therefore considered as having come from multiple distributions. This can be seen clearly in the empirical cdf. The 100 samples from the normal distribution with mean value 6 cause a hump in the cdf centered at 6. These are our fictional “target” detections. The challenge now is to identify these samples and remove them.

Upon further examination of Figure 35, the empirical cdf curve does follow a GPD distribution until it begins flattening out near values of 3. At this point, it intersects the lower bound. Therefore, we can use the lower bound as a threshold for outlier rejection. Any samples in the empirical cdf beyond the lower bound are removed from the data sample. Because the GPD method is a variance reduction

method, it is acceptable to remove some of the non-target samples from the data. This allows us some flexibility in choosing which samples will be used to estimate the new generalized Pareto distribution.

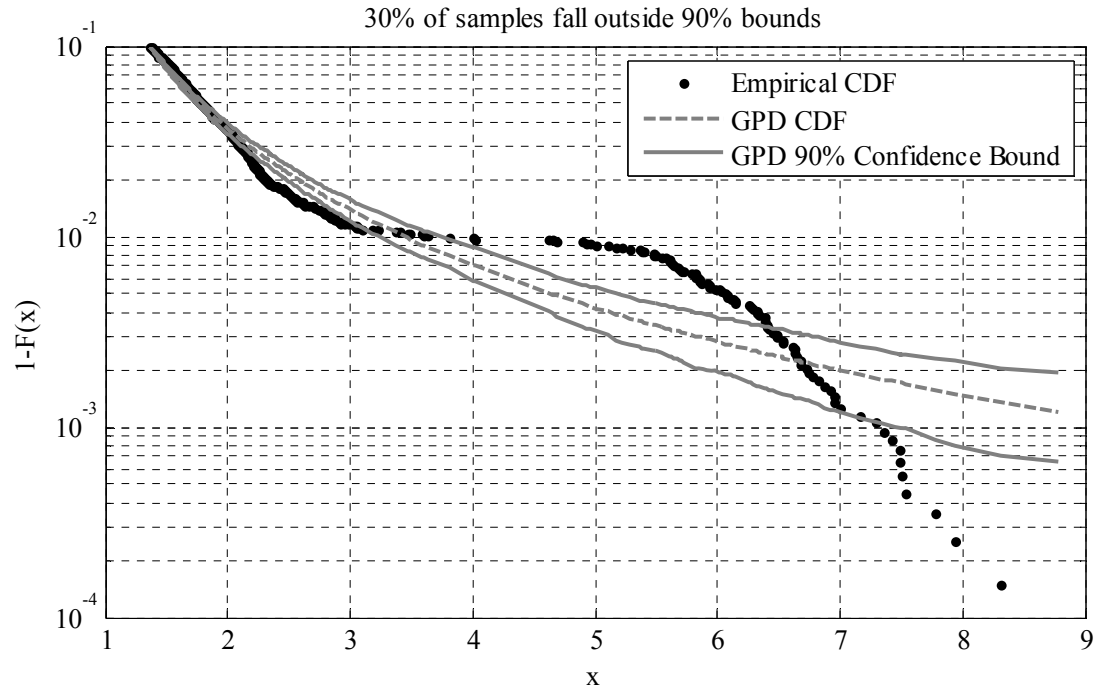


Figure 35: Comparison of the GPD to the Empirical CDF for Example 2

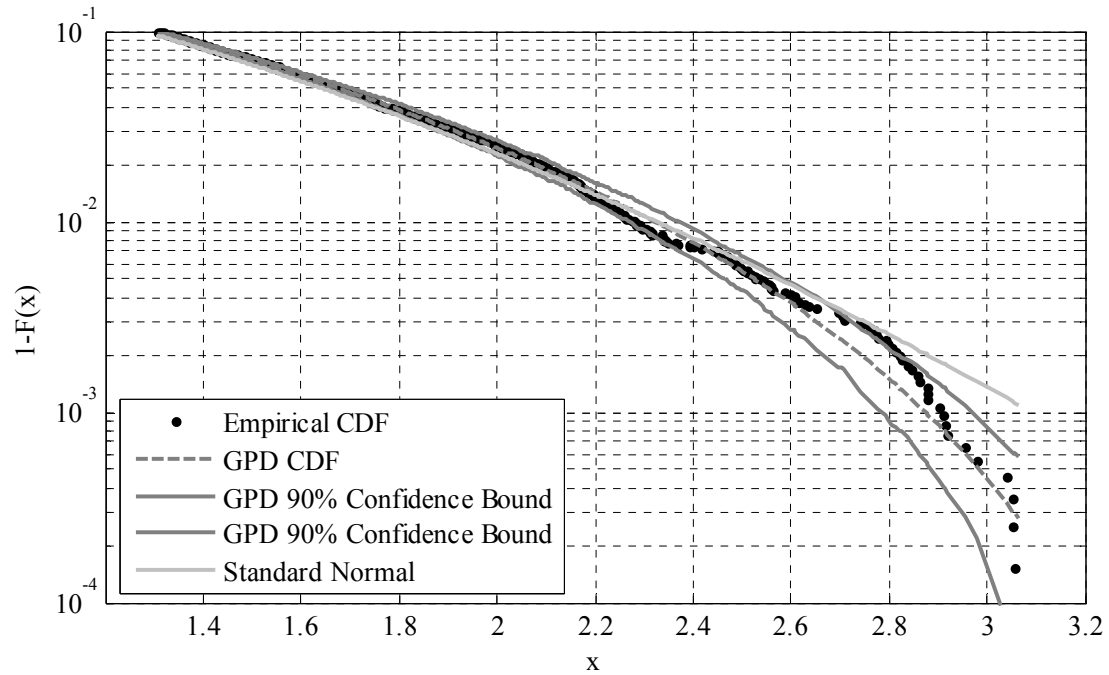


Figure 36: Comparison of Corrected Samples

Using the lower bound to identify the samples to keep, we recalculate the GPD and display the results in Figure 36. Along with the edited empirical cdf and GPD estimates, we include the true cdf of a standard normal distribution. The edited samples now approximate the true normal cdf well –especially at lower samples. The results only diverge at the highest samples and even then, they differ only by 0.0005. This shows that the algorithm can identify samples with “targets”, prune the “target” samples, and then recompute a new tail distribution that is close to the original “background” samples. All of this can be done without any knowledge of the underlying background distribution or knowledge of the target samples. A block diagram of the proposed algorithm is given in Figure 37.

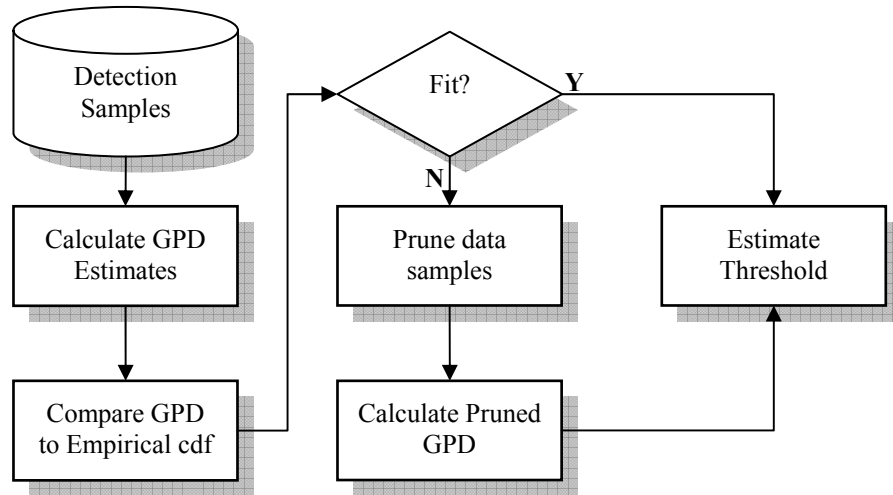


Figure 37: Block Diagram of the EVT Adaptive Threshold Algorithm

6.3. Experimental Results

Our hypothesis is we can detect and eliminate the influence of target samples to adaptively threshold detection results. Not only can we eliminate the influence of the target samples, but by using the generalized Pareto distribution, we can accurately estimate a threshold for a desired false alarm density. To show whether this occurs or

not in practice, we have implemented a number of experiments on both known distributions and on subpixel detector results from real-world hyperspectral imagery. The following sections describe the experimental design philosophy and provide results for two experiments measuring the accuracy of the GPD against known distributions and the ability of the EVT Adaptive Threshold algorithm to determine the thresholds for desired false alarm densities on subpixel detection results.

6.3.1. Experiments with Known Distributions

The first set of experiments shows the ability of the GPD to accurately estimate thresholds on known distributions. We use three distributions for this experiment: the normal distribution, the chi-squared distribution with 169 degrees of freedom, and a beta distribution with parameters 0.5 and 84. The normal distribution was used as a statistical benchmark. The chi-squared distribution was used because it represents the detection output of the well-known RX anomaly detector [84]. Finally, the beta distribution represents the statistical output of the ACE detector introduced in Chapter 5.

Another reason for using these distributions is because they all represent different ranges and limits. The normal distribution is valid for the entire real line. The chi-squared distribution is only valid for non-negative values of the real-line. The most limiting distribution is the beta distribution whose range is restricted between 0 and 1. All of these distributions test the ability of the GPD estimate to adapt to different statistical properties. Again, the GPD knows nothing about the true underlying distribution – only that the various tails of the distributions should converge in probability to a generalized Pareto distribution.

For each of the distributions listed above, a set of experiments was conducted to measure the accuracy and precision of the MC and GPD methods. The experiments were developed to estimate thresholds for false alarm densities at 10^{-2} , 10^{-3} , and 10^{-4} given 1000 samples from the distribution in question. Note that these experiments should task each of the methods by attempting to find thresholds as low as 10^{-4} with only 1000 samples – a threshold beyond the MC method’s abilities. At each of the thresholds, 1000 runs were performed to achieve reasonable measurements of the mean and variance. The results of these experiments are given in Table 11. The table includes estimates for the MC method, the GPD method with clustering threshold of 10%, and the theoretical ideal for each false alarm probability α_0 . For the MC and GPD methods, the table includes the mean with the variance in parentheses for each α_0 .

Table 11: Comparison of MC and GPD on Known Distributions

Distribution	α_0	Ideal	MC	GPD
N(0,1)	10^{-2}	2.326	2.348 (0.016)	2.331 (0.009)
	10^{-3}	3.090	3.233 (0.125)	3.038 (0.053)
	10^{-4}	3.719	3.239 (0.122)	3.517 (0.205)
χ^2_{169}	10^{-2}	187.5	187.8 (5.967)	187.6 (3.556)
	10^{-3}	203.4	206.9 (56.83)	202.3 (24.57)
	10^{-4}	217.0	206.9 (56.48)	213.6 (109.4)
Beta(0.5,84)	10^{-2}	0.0386	0.0393 ($1.1 \cdot 10^{-5}$)	0.0384 ($0.6 \cdot 10^{-5}$)
	10^{-3}	0.0622	0.0675 ($1.6 \cdot 10^{-4}$)	0.0612 ($0.7 \cdot 10^{-4}$)
	10^{-4}	0.0859	0.0685 ($1.7 \cdot 10^{-4}$)	0.0875 ($5.1 \cdot 10^{-4}$)

The results from these experiments demonstrate the theoretical gains of using the GPD method. For all distributions, the GPD method obtains a better estimate of the threshold with nearly half the variance of the MC method. This is expected given the variance reduction benefits of using the generalized Pareto distribution. The GPD is also able to provide an estimate for $\alpha_0 = 10^{-4}$. While the estimate does have some bias, it shows the ability of the GPD to take advantage of its variance reduction property to estimate thresholds beyond that of MC methods.

6.3.2. *Experiments on Subpixel Target Detectors*

The simulated results are good for comparing the GPD method with its MC counterpart, but these experiments do not take into account situations that occur in real HSI data. In these cases, the data may not be necessarily homogeneous and can contain numerous outliers. This is especially true when targets are present in the imagery. To measure the usefulness of the GPD-based EVT adaptive threshold method on such data, we applied it and a number of other well-known techniques to the ACE and HSD detector results from Chapter 5 on Target 2. The ACE results were chosen because ACE has a known output distribution (assuming normal statistics). We chose HSD because the detector's output statistics cannot be easily quantified. Target 2 was chosen because it is not the strongest or weakest target signature and provides a good challenge for the algorithms.

6.3.2.1. *ACE Threshold Results*

For the experiments with the ACE detector, we tested four different algorithms. The parameters for this experiment were set such that the desired false alarm density varied from 10^{-3} to 10^{-5} , P is 1, and L is 169. The first algorithm is based on a theoretical calculation using (57). The second algorithm is a parametric

algorithm based on (57); however, instead of using the theoretical parameters, the parameters are estimated directly from the data. The third algorithm is the MC algorithm. The last algorithm is the proposed EVT method. For the EVT method, we use the clustering threshold of 1% to select the samples for estimation of the GPD parameters. On Images 2, 3, 5, and 6, no targets are present; therefore, the MC method should be ideal. On Images 1 and 4, however, where numerous targets are present in the data, we expect the EVT method to perform best. The results for the ACE detector are in Table 12 through Table 14.

Table 12: Comparison of Threshold Estimates for ACE Results

α_0	Image	Theoretical	Parametric	MC	EVT	Ideal
10^{-3}	1	0.0626	0.0664	0.1136	0.0736	0.0759
	2	0.0626	0.0610	0.0681	0.0695	0.0681
	3	0.0626	0.0668	0.0740	0.0751	0.0740
	4	0.0626	0.0656	0.0970	0.0711	0.0750
	5	0.0626	0.0600	0.0690	0.0707	0.0690
	6	0.0626	0.0682	0.0804	0.0823	0.0804
10^{-4}	1	0.0864	0.0922	0.6428	0.1171	0.1146
	2	0.0864	0.0843	0.1111	0.1063	0.1111
	3	0.0864	0.0923	0.1161	0.1146	0.1161
	4	0.0864	0.0910	0.6951	0.1126	0.1203
	5	0.0864	0.0830	0.1449	0.1123	0.1449
	6	0.0864	0.0944	0.1334	0.1305	0.1334
10^{-5}	1	0.1100	0.1177	0.7644	0.1737	0.1533
	2	0.1100	0.1075	0.2201	0.1515	0.2201
	3	0.1100	0.1175	0.1876	0.1630	0.1876
	4	0.1100	0.1162	0.8396	0.1710	0.1684
	5	0.1100	0.1057	0.2637	0.1669	0.2637
	6	0.1100	0.1201	0.2435	0.1935	0.2435

In each table, there are seven columns. The first column identifies the desired false alarm rate we want to achieve. The second column identifies the image that is being processed. The next four columns give the results for the theoretical, parametric, MC, and EVT methods. The last column presents the ideal results for the

desired false alarm rate. This ideal setting was found using the ground truth information to identify target clusters as described in Chapter 2. These target samples were then removed and the rest of the pixels were ordered by detection score. The MC method was then applied to this reduced set to identify the “ideal” threshold.

Table 13: Comparison of P_d Estimates for ACE Results

α_0	Image	Theoretical	Parametric	MC	EVT	Ideal
10^{-3}	1	1.00	1.00	1.00	1.00	1.00
	2	0.00	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00	0.00
	4	1.00	1.00	1.00	1.00	1.00
	5	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00
10^{-4}	1	1.00	1.00	0.24	1.00	1.00
	2	0.00	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00	0.00
	4	1.00	1.00	0.33	1.00	1.00
	5	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00
10^{-5}	1	1.00	1.00	0.02	0.98	1.00
	2	0.00	0.00	0.00	0.00	0.00
	3	0.00	0.00	0.00	0.00	0.00
	4	1.00	1.00	0.03	1.00	1.00
	5	0.00	0.00	0.00	0.00	0.00
	6	0.00	0.00	0.00	0.00	0.00

The results show the usefulness of the EVT method even when the detector distribution can be assumed. The theoretical calculation using the beta distribution underestimates the thresholds consistently. This leads to false alarm rates that are significantly higher than the desired rates. In the most extreme case of 10^{-5} , the false alarm rate is nearly an order of magnitude greater than the desired rate. While the ACE detector is a CFAR detector, the high false alarm rates occur because the underlying HSI data is rarely normally distributed [103]. This assumption of

normality leads to a mismatch between theory and real HSI data causing the higher false alarms and incorrect thresholds.

The parametric method performs slightly better than the theoretical case. Instead of using the predicted parameters for the beta distribution, the parameters are estimated using the maximum likelihood technique. These estimates do improve the results, but the underlying assumption that the data comes from a normal distribution (thus leading to the beta distribution of the ACE detector) does not match the true distribution of the HSI data. Therefore even with estimated parameters, the parametric method does not perform well.

Table 14: Comparison of False Alarms for ACE Results

α_0	Image	Theoretical	Parametric	MC	EVT	Ideal
10^{-3}	1	256	202	11	120	102
	2	155	180	102	94	102
	3	226	169	102	95	102
	4	216	182	29	125	102
	5	147	181	102	91	102
	6	327	220	102	96	102
10^{-4}	1	50	42	0	9	10
	2	26	29	10	11	10
	3	45	32	10	10	10
	4	55	38	0	13	10
	5	42	45	10	20	10
	6	73	43	10	12	10
10^{-5}	1	13	8	0	0	1
	2	10	10	1	2	1
	3	11	9	1	2	1
	4	16	11	0	0	1
	5	22	23	1	5	1
	6	22	14	1	4	1

The MC estimates are more interesting. As expected, the MC estimates are ideal when no targets are present. If only a few targets are present, the MC estimates will continue to provide good thresholds for larger desired false alarm rates. In these

experiments however, the targets span tens of pixels. While this may not be significant at 10^{-2} , it does affect the P_d and desired false alarm rates at 10^{-3} and below. Because the MC method has no mechanism to identify possible target samples, it degrades as the desired false alarm density becomes small. This has the unfortunate effect of removing target detections first before removing clutter (assuming the detector has done an adequate job of separating the targets from the background). The final result is threshold estimates much higher than the ideal which penalize the P_d .

The EVT method performs well in these experiments. The method was able to isolate the influence of the target signatures in Images 1 and 4 before calculating the threshold. The result is a threshold that is near ideal for false alarm rates of 10^{-3} and 10^{-4} . At these false alarm rates, the method provides P_d and false alarm numbers that are unmatched by any other algorithm when targets are present. At the 10^{-5} false alarm rate, the EVT method begins to diverge from the ideal cases; however, the EVT method still provides thresholds that exceed the ability of the MC method. This is an intriguing result as the EVT method is using less than 10,000 samples to estimate a 10^{-5} desired false alarm rate with good accuracy. When targets are not present, the MC method provides the best results as expected; however, the EVT method provides results that are close to ideal. When considering the EVT method's ability to estimate thresholds close to ideal in the presence or absence of targets, the slight errors in threshold level are acceptable to maintain good performance in all conditions.

6.3.2.2. HSD Threshold Results

For the experiments with the HSD detector, we tested only two algorithms because HSD's use of non-negativity constraints precludes the derivation of a theoretical distribution for the detector. The parameters for this experiment were set

such that the desired false alarm density varied from 10^{-3} to 10^{-5} as in the ACE experiment. The two algorithms tested are the MC and EVT methods. For the EVT method, we use the clustering threshold of 1% to select the samples for estimation of the GPD parameters. On Images 2, 3, 5, and 6, no targets are present; therefore, the MC method should be ideal. On Images 1 and 4, however, where numerous targets are present in the data, we expect the EVT method to perform best. The results for the HSD detector are in Table 15 through Table 17.

Table 15: Comparison of Threshold Estimates for HSD Results

α_0	Image	MC	EVT	Ideal
10^{-3}	1	1.0912	1.0540	1.0529
	2	1.0750	1.0738	1.0750
	3	1.0266	1.0207	1.0266
	4	1.1199	1.0884	1.0934
	5	1.0669	1.0668	1.0669
	6	1.0706	1.0709	1.0706
10^{-4}	1	2.4647	1.1011	1.0912
	2	1.1061	1.1142	1.1061
	3	1.0455	1.0416	1.0455
	4	3.1925	1.1395	1.1491
	5	1.0898	1.0973	1.0898
	6	1.1064	1.1131	1.1064
10^{-5}	1	3.7026	1.1759	1.1124
	2	1.1439	1.1632	1.1439
	3	1.0862	1.0773	1.0862
	4	6.5592	1.2100	1.2148
	5	1.1148	1.1312	1.1148
	6	1.1614	1.1687	1.1614

In each table, there are five columns. The first column identifies the desired false alarm rate we want to achieve. The second column identifies the image that is being processed. The next two columns give the results for the MC method and EVT method. The last column provides the ideal results for the desired false alarm rate. This ideal setting was found using the ground truth information to identify target

clusters as described in Chapter 2. These target samples were then removed and the rest of the pixels were ordered by detection score. The MC method was then applied to this reduced set to identify the “ideal” threshold.

Table 16: Comparison of P_d Estimates for HSD Results

α_0	Image	MC	EVT	Ideal
10^{-3}	1	1.00	1.00	1.00
	2	0.00	0.00	0.00
	3	0.00	0.00	0.00
	4	1.00	1.00	1.00
	5	0.00	0.00	0.00
	6	0.00	0.00	0.00
10^{-4}	1	0.24	1.00	1.00
	2	0.00	0.00	0.00
	3	0.00	0.00	0.00
	4	0.33	1.00	1.00
	5	0.00	0.00	0.00
	6	0.00	0.00	0.00
10^{-5}	1	0.02	0.93	1.00
	2	0.00	0.00	0.00
	3	0.00	0.00	0.00
	4	0.03	1.00	1.00
	5	0.00	0.00	0.00
	6	0.00	0.00	0.00

The results for this experiment support the results found using the ACE detector. In this case, however, the detector statistics are entirely unknown and have to be estimated from the data. As expected, the MC method is ideal when no targets are present in the imagery. Once target detections are present, the MC method performs poorly setting the threshold based on target detection scores. This effect, of course, removes targets while giving improper false alarm rates.

The EVT method is able to isolate the target detections and provide good detection thresholds across all images. In images with targets, the EVT method is able to remove the influence of the target samples and calculate thresholds that are near

ideal. The corresponding P_d and false alarm statistics show good performance across all desired false alarm rates. When targets are not present, the EVT method achieves thresholds close to ideal. Again, the GPD method gives good performance across all images regardless of the detection of targets.

Table 17: Comparison of False Alarm Rates for HSD Results

α_0	Image	MC	EVT	Ideal
10^{-3}	1	10	95	102
	2	102	108	102
	3	102	201	102
	4	27	134	102
	5	102	104	102
	6	102	99	102
10^{-4}	1	0	4	10
	2	10	5	10
	3	10	14	10
	4	0	13	10
	5	10	5	10
	6	10	8	10
10^{-5}	1	0	0	1
	2	1	0	1
	3	1	1	1
	4	0	2	1
	5	1	0	1
	6	1	0	1

6.3.3. Conclusions

The EVT adaptive threshold method was developed to work well across all types of detectors and in the presence of targets. The experimental results demonstrate this ability across two different detectors and at multiple desired false alarm rates – even at rates lower than the number of samples present. Strikingly, the method also excels above the theoretical and parametric methods which are based on the known distribution of the detector (unless the data distribution matches the assumed detector distribution).

The other benefit of the EVT method is the speed of calculation. The method takes less than a second to estimate a threshold given a 256x400 pixel image. The method is scalable to any size image and performs as quickly as any of the other methods. This makes the EVT method accessible to a wide range of target applications beyond subpixel detection.

6.4. *Summary*

We present a new way to adaptively estimate detector thresholds via extreme value theory. The method can be used on any detector type – not just those that are CFAR algorithms. In most real-world cases, the EVT adaptive threshold algorithm can outperform CFAR algorithms due to the inherent mismatch between the model assumptions and the real data. Additionally, the EVT method can work in the presence of target detections while still estimating an accurate threshold for a desired false alarm rate. This ability makes it useful to any number of detection applications – not just physics-based subpixel target detection in HSI data.

Chapter 7: Summary

In this dissertation, we have introduced a number of new algorithms for detection of subpixel targets in hyperspectral imagery. Our approach has been to incorporate the known physics of the problem while taking advantage of statistics to account for the unknown variables. Till this point, we have introduced each algorithm separately to isolate their performance. In this chapter, we introduce how these algorithms work together. From this analysis, we identify new areas of research for subpixel detection. We conclude this chapter by summarizing the new algorithms introduced in this dissertation.

7.1. Cumulative Performance Results

In Chapter 1, we presented a block diagram for subpixel target detection in Figure 2. Using that block diagram, we identified the various areas of subpixel detection where we developed new algorithms. These algorithms were independently updated to identify their performance without the influence of the other algorithms. Unfortunately, this never allowed us to bring all the algorithms together to measure their cumulative performance. This section presents an experiment designed to test the cumulative performance of the proposed algorithms.

Figure 38 presents the proposed subpixel detection system. For target characterization, we use the ARRT algorithm introduced in Chapter 3. For background characterization, we use the IEA algorithm and the SDD algorithms described in Chapter 4. The subpixel detector is the HSD algorithm introduced in Chapter 5. Finally, the EVT Adaptive Threshold Algorithm applies a detection threshold based on a desired false alarm density to the HSD detection scores.

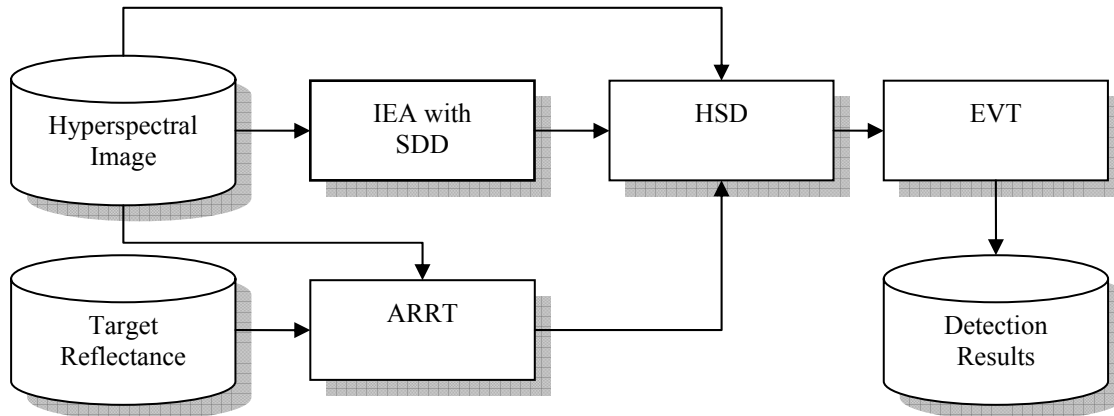


Figure 38: Proposed Subpixel Detection Block Diagram

To show how all of our proposed algorithms work together, we designed an experiment on Target 2. We chose Target 2 because it is not the easiest or hardest target to detect providing a moderate challenge for subpixel detection. We used Images 1 through 6 from Sensor X because these images contain true subpixel targets. The images were left uncalibrated for this experiment to test the ability of the ARRT algorithm to adjust to such conditions. For the target and background reflectance signatures, we used Target 2 and vegetation signatures measured in the field using hand-held spectrometers. No other information was needed to run the system.

The results of the experiment are shown in Figure 39. For reference, we included the best case results for the HSD algorithm operating on Target 2 (as shown in Chapter 5). This best case result assumes the imagery has been vicariously calibrated and target signatures are generated using the MODTRAN algorithm. Additionally, the number of endmembers has been chosen to maximize performance based on ground truth information. This curve represents what a subpixel detector could achieve if all other variables were known.

The dashed gray line is the performance of the HSD algorithm using the EIF background dimension estimate. The EIF method provides consistently good results as shown in Chapter 4. We included this performance curve to show the need for good background dimension estimates even with HSD – a detector partially invariant to the number of background endmembers used.

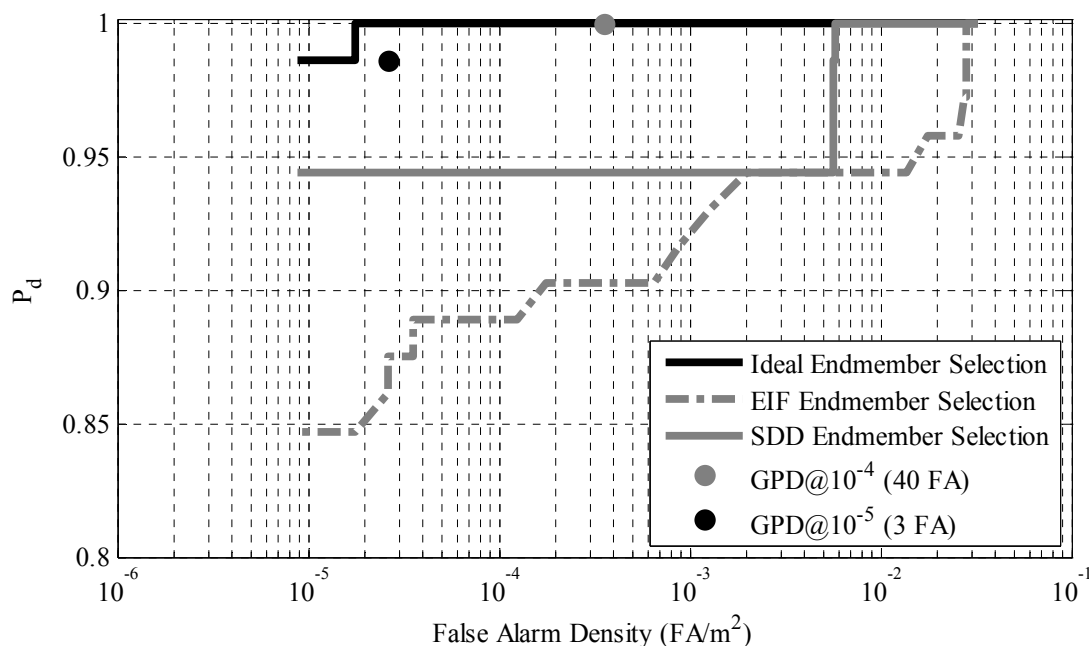


Figure 39: Subpixel Detection System ROC Curves

The solid gray line represents the results of our combined subpixel detector system in Figure 38. This curve shows the system achieves nearly ideal performance. Only two targets are missed at false alarm densities less than 10^{-5} . Even though HSD is partially insensitive to the number of background endmembers chosen, the SDD algorithm is able to produce better results than the EIF algorithm.

Perhaps the most impressive results are the two points calculated by using the EVT Adaptive Threshold Algorithm. The EVT algorithm was applied to the results of the HSD detector (gray line). As noted in Chapter 5, the HSD algorithm sometimes

does not suppress the background into similar ranges of values. The EVT algorithm automatically adapts the threshold for each image taking into account the different background ranges. The result of applying the EVT algorithm provides performance that almost perfectly matches the ideal case. Even though the EVT algorithm is not able to fix the false alarm density exactly, it provides estimates that are very close to the ideal.

The final result is that the proposed combined subpixel detection system is able to obtain performance that is nearly identical to the case where all parameters are known. When one considers the proposed system only uses a target reflectance signature, a reference reflectance signature, and the hyperspectral image without any knowledge of ground truth, the combined performance result is striking. Moreover, the proposed subpixel detection system is able to process each image in less than five minutes making it applicable for near-real time applications.

7.2. Future Work

While this work demonstrates good results for subpixel detection, there are many more interesting topics that spring from the research within this dissertation. Perhaps the most immediate need is improved characterization of target signatures as demonstrated by the subpixel detection results on Target 4. The ARRT and MODTRAN methods both have difficulty handling low reflectance targets. They both produced signatures for Target 4 that underestimated the actual target signature in the SWIR bands. Work should focus on providing better estimates of the upwelled radiance signature using shadow zones as indicated by [80]. These shadow zones can be automatically identified using [1]. Methods can also focus on improved estimates

of the aerosol content of the imagery to help characterize scattering losses at different altitudes.

Estimation of the background dimension remains an active area of research. As shown in Chapter 4, this topic has been only partially treated in the literature. New methods that incorporate target, background, and detector characteristics need to be developed to help improve this area. While our research has produced an improved method to estimate the background dimension, much more could be done.

Another interesting area of research is using the contextual information gained by using physically meaningful endmembers and abundances. For example, when looking for a white automobile, you can remove detections that are not on roads or parking lots. This information can be used to build site models that lead to improved spectral object level change detection (SOLCD) studies [44].

An interesting branch of subpixel detection was proposed by Kwon and Nasrabadi using kernel-based methods [60][61]. The reason for using kernel methods is to project the data into a space that can account for nonlinearities in the data not covered by first and second order moments. They show promising results although their work uses the energy algorithm to estimate the number of background endmembers for the AMSD algorithm [60]. Thus, we cannot identify how well the kernel methods improve detection performance because AMSD performance has been degraded unintentionally.

Nevertheless, the kernel methods open up the possibility of physics-based kernel methods. Just as we created the hybrid detectors by incorporating the known physics of the linear mixing model, we can take the same approach with their kernel

counterparts. For example, research has proposed a new method to extract endmembers based on Support Vector Data Description [6]. This method extracts endmembers in the kernel space that identify the endmembers as the vertices of the enclosing hypersphere. From this work, we developed a Kernel FCLS method to accurately estimate the abundances of those endmembers in the kernel space allowing for the possibility of greater separation between similar spectral signatures [11]. The next step is to modify the Kernel AMSD and Kernel ACE detectors to use the new physics-based kernel parameters. This work will produce a Kernel Hybrid Structured Detector and Kernel Hybrid Unstructured Detector. These algorithms will then be assessed relative to their hybrid counterparts presented in [12]. Other interesting work in kernel methods is the development of algorithms to estimate the kernel parameters – a challenging subject in all kernel methods [92].

While this dissertation focused on the reflective region of the electromagnetic spectrum, hyperspectral sensors have been developed for the Mid-Wave Infrared (MWIR) from 3.0 to 7.0 microns and the Long Wave Infrared (LWIR) from 7.0 to 15.0 microns regions as well. At these wavelengths, emissivity dominates the spectral signature. Emissivity is “the ratio of the emission from [a] material to that of a blackbody at the same temperature” [93]. Therefore, emissivity is a measure of the energy an object emits instead of reflects. Initial work has already been finished applying the hybrid detectors to LWIR sensors [13]. However, target characterization is much more difficult in MWIR and LWIR because temperature has to be accounted for as well as the emissivity [93]. These topics should be pursued however because LWIR sensors provide the opportunity to work in either day or night conditions.

7.3. *Contributions*

In this dissertation, we present a physics-based approach to subpixel detection in hyperspectral imagery. This physics-based approach required the development of new techniques at all levels of subpixel detection from target characterization to threshold estimation. In this section, we summarize the contributions of this thesis:

- We have developed a new target characterization method based on principles of radiative transfer theory and detection theory. Results show this method matches the results by model-based methods, but requires no ancillary data such as weather information, source-target-receiver information, or calibrated sensor responses.
- We have developed a new method to estimate the number of endmembers for subpixel detection applications. We show that the proposed SDD method performs well when compared to the state-of-the-art methods.
- More importantly, we show that for the first time how poor estimates of background dimension lead to significantly reduced subpixel detection performance.
- We created two new physics-based subpixel detectors. The HSD and HUD detectors are the combination of physics-based knowledge to produce physically meaningful parameter estimates and detection theory to account for unknown quantities in the data. Results show these detectors have three advantages: insensitivity to the number of endmembers, improved performance on an image to image basis, and consistent performance across images better than that of known CFAR detectors.

- We developed an adaptive threshold technique based on extreme value theory. This technique is applicable to a wide variety of detectors – not just those that are CFAR. Additionally, the method is able to suppress the influence of target detections to make accurate estimates of the detection threshold without any knowledge of the underlying distribution of the data.

Bibliography

- [1] S.M. Adler-Golden, M.W. Matthew, G.P. Anderson, G.W. Felde, and J.A. Gardner, "Algorithm for de-shadowing spectral imagery," *Proceedings of the SPIE, Imaging Spectrometry VIII*, vol. 4816, pp. 203-210, November 2002.
- [2] H. Akaike, "A new look at the statistical identification model," *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716-723, December 1974.
- [3] G.P. Anderson, B. Pukall, C.L. Allred, L.S. Jeong, M. Hoke, J.H. Chetwynd, S.M. Adler-Golden, A. Berk, L.S. Bernstein, S.C. Richtsmeier, P.K. Acharya, and M.W. Matthew, "FLAASH and MODTRAN4: state-of-the-art atmospheric correction for hyperspectral data," *IEEE Proceedings of the 1999 Aerospace Conference*, vol. 4, pp. 177-181, March 1999.
- [4] E. A. Ashton and A. Schaum, "Algorithms for the detection of sub-pixel targets in multispectral imagery," *Photogramm. Eng. Remote Sensing*, pp. 723-731, July 1998.
- [5] A.A. Balkema and L. de Haan, "Residual life time at great age," *The Annals of Probability*, vol. 2, no. 5, pp. 792-804, October 1974.
- [6] A. Banerjee, P. Burlina, and J. Broadwater, "A machine learning approach for finding hyperspectral endmembers," accepted to *IEEE International Geoscience and Remote Sensing Symposium 2007*, Barcelona, Spain, July 2007.
- [7] C.A. Bateson, G.P. Asner, and C.A. Wessman, "Endmember bundles: A new approach to incorporating endmember variability into spectral mixture analysis," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, pp. 1083-1094, March 2000.
- [8] J. Boardman, "Inversion of high spectral resolution data," *Proc. SPIE*, vol. 1298, pp. 222-233, 1990.
- [9] J.W. Boardman, F.A. Kruse, and R.O. Green, "Mapping target signatures via partial unmixing of AVIRIS data," *Summaries of the VI JPL Airborne Earth Science Workshop*, Pasadena, CA 1995.
- [10] J. Bowles, D. Gillis, and P. Palmadesso, "New improvements in the ORASIS algorithm," *Proceedings of the IEEE Aerospace Conference*, Big Sky, MT, vol. 3, pp. 293-298, March 2000.
- [11] J. Broadwater, A. Banerjee, P. Burlina, and R. Chellappa, "Kernel Fully Constrained Least Squares Abundance Estimates," accepted to *IEEE*

International Geoscience and Remote Sensing Symposium 2007, Barcelona, Spain, July 2007.

- [12] J. Broadwater and R. Chellappa, "Hybrid Detectors for Subpixel Targets," accepted to *IEEE Transactions on Pattern Analysis and Machine Intelligence*, December 2006.
- [13] J. Broadwater and R. Chellappa, "Physics-based detectors applied to long-wave infrared hyperspectral data," *Proceedings of the 2006 Army Science Conference*, November 2006.
- [14] J. Broadwater and R. Chellappa, "An adaptive threshold method for hyperspectral target detection," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. V-1201 – V-1204, May 2006.
- [15] J. Broadwater, R. Meth, and R. Chellappa, "Average relative radiance transform for subpixel detection," *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium 2005*, vol. 5, Seoul, South Korea, pp. 3565-3568, July 2005.
- [16] J. Broadwater, R. Meth, and R. Chellappa, "A Hybrid Algorithm for Subpixel Detection in Hyperspectral Imagery," *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium 2004*, vol. 3, Anchorage, AK, pp. 1601-1604, September 2004.
- [17] J.A. Bucklew, "The blind simulation problem and regenerative processes," *IEEE Transactions on Information Theory*, vol. 44, no. 7, November 1998, pp. 2877-2891.
- [18] G. Casella and R.L. Berger, *Statistical Inference*, Second Ed. Duxbury Press, 2001.
- [19] C-I Chang and Q. Du, "Estimation of Number of Spectrally Distinct Signal Sources in Hyperspectral Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 3, pp. 608-619, March 2004.
- [20] C-I Chang and D.C. Heinz, "Constrained Subpixel Target Detection for Remotely Sensed Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, no. 3, pp. 1144-1159, May 2000.
- [21] R.N. Clark, G.A. Swayze, K.E. Livo, R.F. Kokaly, T.V.V. King, J.B. Dalton, J.S. Vance, B.W. Rockwell, T. Hoefen, and R.R. McDougal, "Surface Reflectance Calibration of Terrestrial Imaging Spectroscopy Data: A Tutorial Using AVIRIS," *Proceedings of the 10th Airborne Earth Science Workshop*, JPL Publication 02-1, 2002.

- [22] D.R. Cox and D. Oakes, *Analysis of Survival Data*, Chapman & Hall, London, 1984.
- [23] R. Davis and S. Resnick, "Tail estimates motivated by extreme value theory," *The Annals of Statistics*, vol. 12, no. 4, pp. 1467-1487, December 1984.
- [24] R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*, Second Edition, John Wiley & Sons, Inc., 2001.
- [25] P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modeling Extremal Events for Insurance and Finance*, Springer, Berlin, 1997.
- [26] W.H. Farrand, R.B. Singer, and E. Merenyi, "Retrieval of apparent surface reflectance from AVIRIS data – a comparison of empirical line, radiative-transfer, and spectral mixture methods," *Remote Sensing of Environment*, Vol. 47, pp. 311-321, 1994.
- [27] M.D. Farrell and R.M. Mersereau, "On the Impact of Covariance Contamination for Adaptive Detection in Hyperspectral Imaging," *IEEE Signal Processing Letters*, vol. 12, no. 9, pp. 649-652, September 2005.
- [28] R.A. Fisher and L.H.C. Tippett, "Limiting forms of the frequency distribution of the largest or smallest member of a sample," *Proceedings of the Cambridge Philosophical Society*, vol. 24, pp. 180-190, 1928.
- [29] G. Foody and D. Cox, "Sub-pixel land cover composition estimation using a linear mixture model and fuzzy membership model and fuzzy membership function," *International Journal of Remote Sensing*, vol. 15, pp. 619-631, 1994.
- [30] M. Frechet, "Sur la loi de probabilité de l'écart maximum," *Ann. Soc. Pol. Math. (Krakow)*, vol. 6, pp. 93-116, 1927.
- [31] J.R. Freemantle, R. Pu, and J.R. Miller, "Calibration of imaging spectrometer data to reflectance using pseudo-invariant features," *Proceedings of the 15th Candian Symposium on Remote Sensing*, pp. 452-455, 1992.
- [32] B.C. Gao and A.F.H. Goetz, "Column atmospheric water vapor and vegetation liquid water retrievals from airborne imaging spectrometer data," *Journal of Geophysical Res. – Atm.*, Vol. 95, pp. 3549-3564, 1990.
- [33] R. Gençay, F. Selçuk, and A. Ulugülyağci, "EVIM: A software package for extreme value analysis in MATLAB," *Studies in Nonlinear Dynamics and Econometrics*, vol. 5, no. 3, pp. 213-244, 2001.
- [34] M. Gianinetto and G. Lechi, "The Development of Superspectral Approaches for the Improvement of Land Cover Classification," *IEEE Transactions on*

Geoscience and Remote Sensing, vol. 42, no. 11, pp. 2670-2679, November 2004.

- [35] B. Gnedenko, "Sur la distribution limite du terme maximum d'une série aléatoire," *Annals of Mathematics*, vol. 44, pp. 423-453, 1943.
- [36] A.A. Green, M. Berman, P. Switzer, and M.D. Craig, "A Transformation for Ordering Multispectral Data in Terms of Image Quality with Implications for Noise Removal," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 26, no. 1, January 1988.
- [37] J.M. Grossman, J. Bowles, D. Haas, J.A. Antoniadis, M.R. Grunes, P. Palmadesso, D. Gillis, K.Y. Tsang, M. Baumbach, M. Daniel, J. Fisher, and I. Triandaf, "Hyperspectral analysis and target detection system for the adaptive spectral reconnaissance program (ASRP)," *SPIE, Algorithms for Multispectral and Hyperspectral Imagery IV*, Orlando, FL, vol. 3372, pp. 2-13, 13-14 April 1998.
- [38] M. Guida, D. Iovino, and M. Longo, "Comparative performance analysis of some extrapolative estimators of probability tails," *IEEE Journal on Selected Areas in Communications*, vol. 6, no. 1, pp. 76-84, January 1988.
- [39] E. Gumbel, *Statistics of Extremes*, Columbia University Press, New York, 1958.
- [40] B. Hapke, *Introduction to the Theory of Reflectance and Emittance Spectroscopy*, Cambridge University Press: Cambridge, UK, 1993.
- [41] J.C. Harsanyi and C-I Chang, "Hyperspectral Image Classification and Dimensionality Reduction: An Orthogonal Subspace Projection Approach," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 32, no. 4, pp. 779-785, July 1994.
- [42] K.H. Haskell and R.J. Hansen, "An algorithm for linear least squares problems with equality and non-negativity constraints generalized," *Math Prog.*, vol. 21, pp. 98-118, 1981.
- [43] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer-Verlag, New York, NY, 2001.
- [44] G.G. Hazel, "Object-level change detection in spectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 3, pp. 553-561, March 2001.
- [45] G. Healey and D. Slater, "Models and Methods for Automated Material identification in Hyperspectral Imagery Acquired Under Unkonwn

- Illumination and Atmospheric Conditions,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 6, pp 2706-2717, November 1999.
- [46] D.C. Heinz and C-I Chang, “Fully Constrained Least Squares Linear Spectral Mixture Analysis Method for Material Quantification in Hyperspectral Imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 3, pp. 529-545, March 2001.
 - [47] V.J. Hodge and J. Austin, “A survey of outlier detection methodologies,” *Artificial Intelligence Review*, vol. 22, pp. 85-126, 2004.
 - [48] A.R. Huete, “A soil-adjusted vegetation index (SAVI),” *Remote Sensing of Environment*, vol. 25, no. 3, pp. 295-309, August 1988.
 - [49] S. Johnson, “The Constrained Signal Detector,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 6, pp. 1326-1337, June 2002.
 - [50] R. Kashyap, “A Bayesian Comparison of Different Classes of Dynamic Models Using Empirical Data,” *IEEE Transactions on Automatic Control*, vol. AC-22, no. 5, pp. 715-727, October 1977.
 - [51] R. Kashyap, “Inconsistency of the AIC Rule for estimating the order of autoregressive models,” *IEEE Transactions on Automatic Control*, vol. AC-25, no. 5, pp. 996-998, October 1980.
 - [52] Y.J. Kaufman, “Solution of the equation of radiative transfer for remote sensing over nonuniform surface reflectivity,” *Journal of Geophysical Research*, vol. 87, no. C6, pp. 4137-4147, 1982.
 - [53] E.J. Kelly, “An Adaptive Detection Algorithm,” *IEEE Transactions in Aerospace and Electronic Systems*, vol. 22, pp. 115-127, March 1986.
 - [54] N. Keshava and J.F. Mustard, "Spectral Unmixing," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 44-57, January 2002.
 - [55] G.R. Kirchhoff and R. Bunsen, “Chemical Analysis by Observation of Spectra,” *Annalen der Physik und der Chemie*, vol. 110, pp. 161-189, 1860.
 - [56] S. Kraut and L.L. Scharf, “The CFAR adaptive sub-space detector is a scale-invariant GLRT,” *IEEE Transactions on Signal Processing*, vol. 47, pp. 2538-2541, Sept. 1999.
 - [57] S. Kraut, L.L. Scharf, and L.T. McWhorter, “Adaptive subspace detectors,” *IEEE Transactions in Signal Processing*, vol. 49, no. 1, pp. 1-16, January 2001.
 - [58] S. Kraut, L.L. Scharf, and R.W. Butler, “The Adaptive Coherence Estimator: A Uniformly Most-Powerful-Invariant Adaptive Detection Statistic,” *IEEE*

Transactions on Signal Processing, vol. 53, no. 2, pp. 427-438, February 2005.

- [59] F.A. Kruse, "Use of Airborne Imaging Spectrometer Data to Map Minerals Associated with Hydrothermally Altered Rocks in the Northern Grapevine Mountains, Nevada, and California," *Remote Sensing of the Environment*, vol. 24, no. 1, pp. 31-51, February 1988.
- [60] H. Kwon and N.M. Nasrabadi, "Kernel matched subspace detectors for hyperspectral target detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 2, pp. 178-194, February 2006.
- [61] H. Kwon and N.M. Nasrabadi, "Kernel adaptive subspace detector for hyperspectral imagery," *IEEE Geoscience and Remote Sensing Letters*, vol. 3, no. 2, pp. 271-275, April 2006.
- [62] J.C. Lagarias, J.A. Reeds, M.H. Wright, and P.E. Wright, "Convergence properties of the Nelder-Mead simplex method in low dimensions," *SIAM Journal of Optimization*, vol. 9, no. 1, pp. 112-147, 1998.
- [63] D. Landgrebe, "The evolution of landsat data analysis," *Photogrammetric Engineering and Remote Sensing*, vol. LXIII, no. 7, pp 859-867, July 1997.
- [64] D. Landgrebe, "Multispectral Land Sensing: Where From, Where to?," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 43, no. 3, pp. 414-421, March 2005.
- [65] K. Lee, "A Subpixel Scale Target Detection Algorithm for Hyperspectral Imagery," Ph.D. Dissertation, Rochester Institute of Technology, Center for Imaging Science, 2002.
- [66] M. Lewis, V. Jooste, and A.A. de Gasparis, "Discrimination of Arid Vegetation with Airborne Multispectral Scanner Hyperspectral Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 7, pp. 1471-1479, July 2001.
- [67] Lord Rayleigh (J.W. Strutt), "On the Light from the Sky, Its Polarization and Colour," *Philosophical Magazine*, vol. 41., pp. 107-120, 274-279, 1871.
- [68] E.R. Malinowski, "Theory of error in factor analysis," *Analytical Chemistry*, vol. 49, no. 4, pp. 606-612, 1977.
- [69] D. Manolakis, "Realistic matched filter performance prediction for hyperspectral target detection," *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium, 2004*, vol. 2, pp. 953-955, September 2004.

- [70] D. Manolakis and G. Shaw, "Detection Algorithms for Hyperspectral Imaging Applications," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 29-43, January 2002.
- [71] D. Manolakis, C. Siracusa, and G. Shaw, "Hyperspectral Subpixel Target Detection Using the Linear Mixing Model," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 7, pp. 1392-1409, July 2001.
- [72] R. Marion, R. Michel, and C. Faye, "Measuring Trace Gases in Plumes from Hyperspectral Remotely Sensed Data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 4, pp. 854-864, April 2004.
- [73] S.S. McArdle, J.R. Miller, and J.R. Freemantle, "Airborne acquisition under clouds: preliminary comparisons with clear-sky scene radiance and reflectance imagery," *Proceedings of the 15th Canadian Symposium on Remote Sensing*, pp. 446-449, 1992.
- [74] L.B. Milstein, D.L. Schilling, and J.K. Wolf, "Robust detection using extreme-value theory," *IEEE Transactions on Information Theory*, vol. IT-15, no.3, pp. 370-375, May 1969.
- [75] P. K. Murphy and M. A. Kolodner, "Classification of Hyperspectral Data for Landuse Applications, Final Report", JHU/APL Report, SRM-01-063, February 8, 2002.
- [76] J.F. Mustard and C.M. Pieters, "Photometric phase functions of common geologic minerals and applications to quantitative analysis of mineral mixture reflectance spectra," *J. Geophys. Res.*, vol. 94, pp. 13619-13634, 1989.
- [77] R.A. Neville, K. Staenz, T. Szeredi, J. Lefebvre, and P. Hauff, "Automatic endmember extraction from hyperspectral data for mineral exploration," in *4th Int. Airborne Remote Sensing Conf. Exhibition/21st Canadian Symposium on Remote Sensing*, Ottawa, Ontario, Canada, 21-24, pp. 891-896, June 1999.
- [78] J. Oberg, "High-tech sensor in the shuttle search," *MSNBC News*, <http://www.msnbc.msn.com/id/3077570/>, Houston, March 30, 2003
- [79] J. Pickands, "Statistical inference using extreme order statistics," *The Annals of Statistics*, vol. 3, no. 1, pp. 119-131, 1975.
- [80] K.R. Piech and J.E. Walker, "Interpretation of Soils," *Photogrammetric Engineering and Remote Sensing*, vol. 40, pp. 87-94, 1974.
- [81] A. Plaza, P. Martínez, R. Pérez, and J. Plaza, "Spatial/spectral endmember extraction by multidimensional morphological operations," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, pp. 2025-2041, September 2002.

- [82] A. Plaza, P. Martínez, R. Pérez, and J. Plaza, "A quantitative and comparative analysis of endmember extraction algorithms from hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 3, pp. 650-663, March 2004.
- [83] R. Price, C.D. Anger, and S. Mah, "Preliminary evaluation of *casi* preprocessing techniques," *Proceedings of the 17th Canadian Symposium on Remote Sensing*, pp. 694-697, 1995.
- [84] I.S. Reed and X. Yu, "Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 10, pp. 1760-1770, October 1990
- [85] H. Ren and C-I Chang, "Automatic Spectral Target Recognition in Hyperspectral Imagery," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1232-1249, October 2003.
- [86] H. Ren, Q. Du, and J. Jensen, "Constrained Weighted Least Squares Approaches for Target Detection and Classification in Hyperspectral Imagery," *Proceedings of the IEEE International Geoscience and Remote Sensing Symposium 2002*, Toronto, Canada, pp. 3426-3428, June 2002.
- [87] J. Rissanen, "A universal prior for integers and estimation by minimum description length," *Annals of Statistics*, vol. 11, pp. 416-431, 1983.
- [88] D.A. Roberts, M.O. Smith, and J.B. Adams, "Green Vegetation, Nonphotosynthetic Vegetation, and Soils in AVIRIS Data," *Remote Sensing of the Environment*, vol. 44, pp. 3039-3054, 1993.
- [89] S.J. Roberts, "Novelty detection using extreme value statistics," *IEE Proceedings of Vision, Image, and Signal Processing*, vol. 146, no. 3, pp. 124-129, June 1999.
- [90] J.W. Rouse, R.H. Haas, J.A. Schell, and D.W. Deering, "Monitoring vegetation systems in the great plains with ERTS," *Third ERTS Symposium*, NASA SP-351, pp. 309-317, 1973.
- [91] J.S. Sadowsky and J.A. Bucklew, "On large deviations theory and asymptotically efficient Monte Carlo simulation," *IEEE Transactions on Information Theory*, vol. 36, issue 3, May 1990, pp. 579-588.
- [92] B. Scholkopf and A.J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, 2002.
- [93] J.R. Schott, *Remote Sensing: The Image Chain Approach*, Oxford, U.K.: Oxford Univ. Press, 1997.

- [94] J.R. Schott, K. Lee, R. Raqueno, and G. Hoffman, "Use of Physics Based Models in Hyperspectral Image Exploitation," *IEEE Proc. of the 31st Applied Imagery Pattern Recognition Workshop (AIPR '02)*, pp. 36-42, October 2002.
- [95] R.A. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*, San Diego, CA: Academic, 1997.
- [96] G. Schwartz, "Estimating the dimension of a model," *The Annals of Statistics*, vol. 5, no. 2, pp. 461-464, 1978.
- [97] S. Seager, E.L. Turner, J. Schafer, and E.B. Ford, "Vegetation's Red Edge: A Possible Spectroscopic Biosignature of Extraterrestrial Plants," *Astrobiology*, vol. 5, no. 3, pp. 372-390, 2005.
- [98] J. J. Settle and N. A. Drake, "Linear mixing and estimation of ground cover proportions," *Int. J. Remote Sensing*, vol. 14, no. 6, pp. 1159-1177, 1993.
- [99] D. Siegmund, "Importance sampling in the Monte Carlo study of sequential tests," *Annals of Statistics*, vol. 4, pp. 673-684, 1976.
- [100] R. Smith, "Extreme value analysis of environmental time series: an application to trend detection in ground-level ozone," *Statistical Science*, vol. 4, pp. 367-393, 1989.
- [101] R. Srinivasan, *Importance Sampling: Applications in Communications and Detection*, Springer-Verlag, Berlin, Germany, 2002.
- [102] R. Srinivasan, "Simulation of CFAR detection algorithms for arbitrary clutter distributions," *IEE Proceedings, Radar, Sonar, and Navigation*, vol. 147, no. 6, pp. 31-40, December 2000.
- [103] D.W. Stein, S.G. Beaven, L.E. Hoff, E.M. Winter, A.P. Shaum, and A.D. Stocker, "Anomaly Detection from Hyperspectral Imagery," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 58-69, January 2002.
- [104] A.D. Stocker and P. Schaum, "Applications of stochastic mixing models to hyperspectral detection problems," *Proceedings of SPIE, Algorithms for Multispectral and Hyperspectral Imagery III*, vol. 3071, Orlando, FL, pp. 47-60, April 1997.
- [105] A. Stratton, *Electromagnetic Theory*, McGraw-Hill, New York, NY, 1941.
- [106] S. Subramanian and N. Gat, "Subpixel Object Detection Using Hyperspectral Imaging for Search and Rescue Operations," *Proc. SPIE for Automatic Target Recognition VIII*, vol. 3371, pp. 216-225, November 1998.

- [107] G. Swayze, R.N. Clark, F. Kruse, S. Sutley, and A. Gallagher, "Ground-truthing AVIRIS Mineral Mapping at Cuprite, Nevada," *Summaries of the Third Annual JPL Airborne Geoscience Workshop*, vol. 1, pp. 47-49, 1992.
- [108] D.M.J. Tax and R.P.W. Duin, "Support vector domain description," *Pattern Recognition Letters*, vol. 20, pp. 1191-1199, 1999.
- [109] B. Thai and G. Healey, "Invariant Subpixel Material Detection in Hyperspectral Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 3, pp. 599-608, March 2002.
- [110] S. Tompkins, J.F. Mustard, C.M. Pieters, and D.W. Forsyth, "Optimization of endmembers for spectral mixture analysis," *Remote Sensing of Environment*, vol. 59, pp. 472-489, 1997.
- [111] G. Vane, R.O. Green, T.G. Chrien, H.T. Enmark, E.G. Hansen, and W.M. Porter, "The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sensing of the Environment*, vol. 44, pp. 127-143, 1993.
- [112] R. von Mises, "La distribution de la plus grande de n valeurs," *Rev. Math. Union Interbalkan*, vol. 1, pp. 141-160, 1936.
- [113] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 2, pp. 387-392, April 1985.
- [114] S.B. Weinstein, "Theory and application of some classical and generalized asymptotic distributions of extreme values," *IEEE Transactions on Information Theory*, vol. IT-19, no. 2, pp. 148-154, March 1973.
- [115] M.E. Winter, "Fast autonomous spectral endmember determination in hyperspectral data," *Proc. 13th Int. Conf. Applied Geologic Remote Sensing*, Vol. II, Vancouver, BC, Canada, pp. 337-344, 1999.
- [116] M.E. Winter, "N-FINDR: An algorithm for fast autonomous spectral end-member determination in hyperspectral data," *Proceedings of SPIE*, vol. 3753, pp. 266-275, 1999.